Risk in a Data-Rich Model*

Dario Caldara[1], Haroon Mumtaz[2], and Molin Zhong[1]

[1]Federal Reserve Board

[2]Queen Mary University of London

March 9, 2026

Abstract

We characterize asymmetric tail risk across over one hundred U.S. macroeconomic and financial variables using a dynamic factor model with stochastic volatility. The model unifies growth-at-risk, inflation-at-risk, and sectoral heterogeneity through common factors whose volatility responds endogenously to shocks, combined with heterogeneous factor loadings. We find that asymmetric tail risk is pervasive and heterogeneous: some sectors exhibit severe asymmetry while others show minimal asymmetry, with variation across activity, price, and financial variables. The framework disentangles supply- and demand-driven tail risk dynamics, revealing how the balance of risks shifts across episodes, and identifies where vulnerabilities concentrate across the economy.

JEL Classification: C11; C32; C38; E32; E44.

Keywords: Dynamic Factor Model; Tail Risk; Stochastic Volatility; Leverage Effect; Growth-at-Risk; Sectoral Heterogeneity.

*Dario Caldara (dario.caldara@frb.gov); Haroon Mumtaz (h.mumtaz@qmul.ac.uk); Molin Zhong (molin.zhong@frb.gov). We thank for helpful comments Todd Clark, Domenico Giannone, Francesca Loria, Matteo Luciani, Michele Modugno, seminar participants at the Adam Smith Business School, the Federal Reserve Board, 2024 BSE Summer Forum on "Advances in Structural Shocks Identification," 2024 workshop on "The Economics of Risk" Turin, 2024 Kansas Econometrics Workshop, 2024 So.Fi.E. workshop on "Monitoring and Forecasting Macroeconomic and Financial Risk", 5th DC-VA-MD Econometrics Workshop, 2024 System Econometrics Conference, 2025 Nowcasting Workshop at PSE, 2025 IAAE Annual Conference, 2025 NBER-NSF Time Series, and 3rd NLMacro Workshop, Lancaster University. We also would like to thank Lilliana Wells, Sofie Grouws, and Charlotte Singer for outstanding research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

1 Introduction

Some sectors of the economy contract sharply during recessions while others remain stable. During the 2008 financial crisis, for instance, motor vehicle production plummeted 40 percent and durable goods consumption declined severely, while electricity generation dropped by less than 5 percent and food consumption remained relatively steady. All sectors experienced the same macroeconomic environment: the same aggregate shocks, the same monetary policy response, the same measures of financial stress. Yet their tail risk profiles differed dramatically. This heterogeneity is not isolated to the global financial crisis (GFC), nor to activity. Across episodes, some sectoral activity, price, and financial variables consistently exhibit extreme risks while others remain relatively stable, with some sectors shifting from symmetric to highly asymmetric risk depending on the nature of shocks. Standard measures of aggregate uncertainty rise uniformly during crises, offering no explanation for this variation. Where does this heterogeneity come from? Is it idiosyncratic noise, or does it reflect systematic propagation of common shocks?

Using a dynamic factor model with endogenous stochastic volatility applied to over one hundred U.S. variables, we show that the heterogeneous tail risk is systematic, arising from common macroeconomic dynamics that propagate asymmetrically across the economy. The core mechanism is an extension of the leverage effect documented in asset markets (Black, 1976; Christie, 1982) and aggregate volatility (Ludvigson et al., 2021): shocks affect factor levels, and factor volatility responds endogenously. For real activity, adverse shocks depress levels while elevating volatility, amplifying downside risk. For prices and financial conditions, shocks can increase both levels and volatility—rising inflation or spreads accompanied by rising uncertainty—generating upside risk. Heterogeneity across sectors reflects differential sensitivity to these dynamics. Cyclical sectors sensitive to financial conditions face severe downside risk amplification. Price-sensitive goods exhibit inflation-at-risk during supply disruptions. Stable sectors remain relatively insulated across episodes.

These dynamics unify three empirical phenomena previously studied in isolation. First, the growth-at-risk behavior documented by Adrian et al. (2019), where GDP growth's lower tail is more volatile than its upper tail, reflects these dynamics operating in real activity variables. Second, inflation-at-risk and financial-stress-at-risk observed across prices and financial variables reflect the same mechanism with opposite sign: rising inflation or spreads accompanied by rising uncertainty generate upside rather than downside risk. Third, the substantial heterogeneity across sectors, from extreme asymmetry in cyclical industries to near-symmetry in stable sectors, reflects how different variables respond to

common underlying shocks. Rather than requiring separate models for growth-at-risk, inflation-at-risk, and sectoral heterogeneity, a unified framework with common macroeconomic fluctuations simultaneously generates the full spectrum of tail risk behavior observed across over one hundred variables.

Our methodological contribution is a dynamic factor model that formalizes these dynamics through three interconnected mechanisms that allow factor levels and volatilities to interact. First, past volatility directly affects current factor dynamics, creating channels from uncertainty to real outcomes. Second, factor movements influence volatility evolution, making uncertainty endogenous to economic conditions rather than exogenous shocks. Third, factor levels and their volatilities are contemporaneously correlated, so disturbances simultaneously affect both. Combined with flexibly estimated factor loadings, these features enable the model to generate the rich heterogeneity in tail risk documented above while maintaining a parsimonious low-dimensional factor structure. The model produces full predictive distributions for all variables, not just point forecasts or volatility estimates, allowing us to characterize time-varying asymmetries in upside and downside risks across the entire macroeconomy.

We estimate the model using 116 monthly U.S. macroeconomic and financial time series from 1973 to 2023, covering major crises including the Great Inflation, the GFC, and the pandemic. The dataset spans real activity, prices, credit, and financial conditions with substantial sectoral disaggregation: industrial production at the three-digit NAICS level, personal consumption expenditures by product category, and equity returns by industry portfolio. This granularity is essential for documenting the heterogeneity in tail risk and linking it to observable sectoral characteristics.

We establish four empirical findings. First, we construct three directional indices—for growth, financial conditions, and inflation—that synthesize tail risk information across many variables. This aggregation approach is more robust than examining individual variables in isolation: the indices are less sensitive to idiosyncratic shocks affecting any single series and provide stable measures of economy-wide tail risk. The indices reveal meaningful time variation consistent with regime shifts. Growth risk tilts strongly toward adverse events during financial crises, when falling activity and rising uncertainty reinforce each other. Inflation risk shifts direction across episodes: upward during supply-shock periods (1970s, 2021-22), downward during demand-driven downturns (2008-09). These patterns align with narrative accounts of different shocks driving these historical episodes.

Second, we examine heterogeneity at the sectoral level. Cyclical, capital-intensive industries like construction materials exhibit extreme downside asymmetry, while regulated

utilities show modest, nearly symmetric risk. Durable goods consumption displays strong growth-at-risk, while nondurable goods like food and necessities remain relatively stable. We validate that these patterns capture economic fundamentals rather than statistical artifacts by linking industrial production risk to Census Bureau data on job flows and establishment dynamics. Sectors with more left-skewed job creation and net establishment entry rates also exhibit more pronounced tail risk asymmetries.

Third, we analyze the estimated factor structure. The latent factors capture financial conditions, consumption, credit, and inflation dynamics, spiking during crises. Mean-volatility correlations vary both within and across factors. Within factors, credit shows negative correlations (declining credit raises volatility), while financial and inflation factors show positive correlations. Across factors, stress in one domain amplifies uncertainty in others. Combined with heterogeneous factor loadings—for instance construction materials and durables load heavily on financial factors, while utilities and food show weak loadings—these correlation patterns generate the heterogeneous tail risk documented above.

Fourth, we validate the model's mechanisms through counterfactual analysis. A financial shock during the GFC generates severe downside risk for credit-sensitive sectors like construction materials and durables, while stable sectors like utilities remain unaffected, closely matching observed 2008-09 patterns. The model also captures regime-dependent risk dynamics: comparing the Great Inflation (1978-82) to the GFC reveals that the same factor structure produces both inflation-at-risk and growth-at-risk when supply shocks dominate, while generating primarily growth-at-risk when demand shocks dominate, consistent with state-dependent propagation in models with occasionally binding constraints.

We assess model fit and validate the sources of heterogeneity through three exercises. First, we show that common factors explain substantial variation in the data—around 80 percent for some key series. Second, we confirm that heterogeneity in tail risk arises from factor loadings rather than idiosyncratic noise: shutting off idiosyncratic shocks preserves the structure of risk across growth, inflation, and financial variables, but randomly reshuffling factor loadings destroys it. Third, comparing conditional distributions to those from quantile regressions (Adrian et al., 2019) across in-sample and out-of-sample forecasts reveals complementary strengths: our factor model outperforms for the tails of real activity and inflation variables, while quantile regressions conditioning directly on financial conditions excel for financial variables.

Our work contributes to three literatures. First, we extend research on macroeconomic uncertainty (Bloom, 2009; Jurado et al., 2015; Carriero et al., 2018) by shifting focus from

symmetric volatility to directional tail risks. Existing uncertainty measures tell us when dispersion is elevated but not whether risks tilt toward adverse or favorable outcomes—a distinction that is critical for assessing the balance of risks.

Second, we contribute to the growth-at-risk literature initiated by Adrian et al. (2019) and extended by López-Salido and Loria (2024), Kiley (2022), and Amburgey and McCracken (2025), which documents asymmetric tail behavior for individual variables using quantile regressions, and Opschoor et al. (2025), which does so across manufacturing industries. We show this phenomenon is pervasive across the economy and provide a unifying mechanism: endogenous mean-volatility interactions in common factors alongside heterogeneous loadings. Growth-at-risk reflects negative loadings on inflation and financial factors, which exhibit positive mean-volatility correlations. Inflation-at-risk and financial-stress-at-risk reflect positive loadings on these same factors. Sectoral heterogeneity reflects heterogeneous factor loadings interacting with these correlations. Rather than separate models for each phenomenon, a small number of common factors with endogenous stochastic volatility simultaneously generates the full spectrum of tail risk behavior across over one hundred variables, moving the literature from documenting patterns to understanding their origins.

Third, we extend dynamic factor models (Stock and Watson, 2002; Bernanke et al., 2005; Stock and Watson, 2016) by incorporating endogenous stochastic volatility. Canova and Ravn (2000) and Canova and Ciccarelli (2013) demonstrate the importance of modeling sectoral and cross-country heterogeneity in panel frameworks, yet typically assume constant volatility. While Del Negro and Otrok (2008) and Mumtaz and Surico (2012) introduce time-varying volatility in factor models, it evolves exogenously in their frameworks—volatility shocks do not affect factor levels, nor do factor movements influence volatility. Carriero et al. (2018) and Davidson et al. (2025) model common volatility with mean effects but specify it as a separate latent factor rather than allowing factors' own volatility to respond to their movements. Gorodnichenko and Ng (2017) demonstrate the importance of separately identifying level and volatility factors in large panels through indirect extraction methods. We extend this insight by modeling factor volatility explicitly through stochastic processes that respond endogenously to economic conditions. Our approach differs by allowing factor movements to directly influence their own volatility and permitting correlation between level and volatility shocks. These features enable our model to generate the asymmetric tail risk patterns we document across over one hundred variables while preserving interpretable factor loadings that vary meaningfully across sectors.

Beyond these methodological and empirical contributions, the framework provides empirical targets for validating dynamic models with nonlinear mechanisms. Recent work on occasionally binding constraints (Guerrieri and Lorenzoni, 2017) and financial frictions predicts that binding episodes should generate asymmetric tail dynamics: severe downside risk when constraints tighten, modest upside risk when they relax. Models incorporating stochastic volatility (Fernández-Villaverde et al., 2011) predict asymmetric real effects of uncertainty shocks. Multi-sector models (Canova and Ravn, 2000) further suggest these effects should vary systematically across industries. Our framework offers not only distributional targets but also interpretable factor dynamics and counterfactual capabilities, illustrated through an historical decomposition of risks during the Great Inflation versus the GFC, that can discipline specific transmission mechanisms, extending validation beyond traditional likelihood-based and moment-matching methods (Fernández-Villaverde and Rubio-Ramírez, 2007; Canova and Sala, 2009).

2 The Model

In this section, we present the dynamic factor model (DFM) with stochastic volatility that forms the foundation of our analysis. We begin by presenting the model specification and then explain why standard DFMs cannot capture the risk dynamics we document. We show how our modeling choices—particularly the correlation between factor levels and their volatilities—generate asymmetric tail risk behavior across many variables. Finally, we present the estimation algorithm and details about the data and empirical specification.

2.1 A DFM with Endogenous Stochastic Volatility

The DFM is represented in a state-space form, where we specify a small number of latent factors F_t that drive the common dynamics of n observed variables X_t, where n is a large number (116 in our application). The observation equation links the observed variables to the latent factors, while the transition equation governs the evolution of these factors over time, incorporating both level and volatility shocks.

The observation equation is written as:

$$X_t = BF_t + v_t, \tag{1}$$

where X_t is an $n \times 1$ vector of observed variables, F_t is an $N \times 1$ vector of latent factors, and v_t represents idiosyncratic components. The factor loadings are captured in

the B matrix, and the idiosyncratic shocks follow an autoregressive (AR) process:

$$v_{it} = \sum_{l=1}^{L} \rho_i v_{it-l} + u_{it}, \quad u_{it} \sim N(0, R_t), \tag{2}$$

where the diagonal covariance matrix $R_t = \text{diag}(e^{r_t})$ models time-varying idiosyncratic volatility, with r_t evolving as independent random walks. We assume idiosyncratic volatilities evolve independently (i.e., R_t is diagonal) to isolate the common risk dynamics driven by the factors. This structure ensures that individual shocks exhibit heteroscedasticity without direct interactions across variables or with the common factors.

The transition equation specifies the evolution of latent factors and incorporates both autoregressive dependencies and volatility effects:

$$F_t = c + \sum_{j=1}^{P} \beta_j F_{t-j} + \sum_{k=1}^{K} b_k \tilde{h}_{t-k} + H_t^{1/2} e_t. \tag{3}$$

The term $\sum_{k=1}^{K} b_k \tilde{h}_{t-k}$ captures the in-mean effect of past volatility on factor dynamics. This mechanism allows elevated past uncertainty to directly affect current economic conditions. For instance, heightened financial volatility may depress real activity or amplify risk premia, consistent with uncertainty-driven business cycle models.

The stochastic volatility process H_t evolves as a VAR:

$$\tilde{h}_{t+1} = \alpha + \theta \tilde{h}_t + \sum_{j=1}^{Q} d_j F_{t-j} + S^{1/2} \eta_t, \tag{4}$$

where $\tilde{h}_t = [h_{1t}, h_{2t}, \ldots, h_{N,t}]$ is a vector of log stochastic volatilities, and $H_t = \text{diag}\left(\exp(\tilde{h}_t)\right)$. Crucially, equation 4 allows factor movements to directly influence volatility evolution, introducing endogeneity between the level and volatility of the factors.

The disturbances $\varepsilon_t = [\eta_t, e_t]$ follow a normal distribution with a time-invariant covariance matrix:

$$\Sigma = \begin{pmatrix} \Sigma_\eta & \Sigma'_{\eta e} \\ \Sigma_{\eta e} & \Sigma_e \end{pmatrix}, \tag{5}$$

where we normalize the diagonal elements of Σ to equal 1 for identification. This normalization, combined with restrictions on the factor loading matrix B described below, ensures the model is identified while allowing flexible correlation patterns between level and volatility shocks. The covariance matrix of the reduced form residuals in equations 3 and 4 is time-varying due to stochastic volatility:

$$\Omega_t = \begin{pmatrix} S^{1/2} & 0 \\ 0 & H_t^{1/2} \end{pmatrix} \begin{pmatrix} \Sigma_\eta & \Sigma'_{\eta e} \\ \Sigma_{\eta e} & \Sigma_e \end{pmatrix} \begin{pmatrix} S^{1/2} & 0 \\ 0 & H_t^{1/2} \end{pmatrix}'. \tag{6}$$

This formulation captures dynamic risk structures by allowing factor shocks and their volatilities to interact over time.

2.2 How the Model Generates Asymmetric Risk Dynamics

We begin by clarifying what we mean by uncertainty, risk, and risk asymmetry, as these concepts are central to our analysis. For any variable in our dataset, such as GDP growth or inflation, we construct the conditional distribution of future outcomes h periods ahead, denoted $p(X_{i,t+h}|X^t)$. This distribution summarizes the full range of possible realizations given current information.

We define *uncertainty* as the dispersion or width of this conditional distribution, measured for instance by its interquartile range or the standard deviation. We define *downside risk* as the 5th percentile of the distribution and *upside risk* as the 95th percentile. These tail quantiles capture the extreme adverse and favorable outcomes that might occur.

Crucially, we are interested in how these measures evolve over time. *Risk asymmetry* occurs when the volatility of the lower tail differs from the volatility of the upper tail across time. For example, if the 5th percentile exhibits substantial variation during recessions while the 95th percentile remains relatively stable, we say the variable exhibits asymmetric downside risk, or "growth-at-risk" behavior (Adrian et al., 2019). This time variation is policy-relevant: central banks and fiscal authorities need to understand not just the average outlook, but how the balance of risks shifts across the business cycle.

The key challenge is generating these time-varying, asymmetric risk dynamics in a model that spans many variables simultaneously. Standard approaches, as we discuss below, cannot capture this behavior.

How Our Model Works: An Illustrative Example

Our model generates asymmetric tail risk through the *correlation between factor levels and their volatilities*. When a factor declines and its volatility increases simultaneously, the lower tail of the conditional distribution shifts substantially leftward (both forces work in the same direction), while the upper tail shifts by less (the forces partially offset). This mechanism, introduced by Adrian et al. (2019) for univariate models and extended to VARs by Caldara et al. (2021), is what enables our model to capture rich risk dynamics across many variables.

Figure 1 illustrates this mechanism using a simplified two-factor model.[1] The top row shows conditional distributions for two factors with opposite mean-volatility correlations. Factor 1 exhibits positive correlation: when it rises, uncertainty increases, creating a long right tail with elevated upside risk. Factor 2 exhibits negative correlation: when it falls, uncertainty increases, creating a long left tail with elevated downside risk. The vertical lines mark the 5% and 95% quantiles, showing how the tails shift asymmetrically.

The bottom row demonstrates the key insight: observable variables inherit different risk patterns through the factor loading matrix B. Variable X_a loads negatively on factors with positive mean-volatility correlation, producing growth-at-risk dynamics (volatile left tail). Variable X_b loads on both factors, yielding symmetric risk. Variable X_c loads positively on factors with positive correlation, generating inflation-at-risk dynamics (volatile right tail). Thus, a small number of factors with distinct correlation patterns can generate diverse risk behaviors across many variables through linear combinations.

This parsimony is the key advantage of our approach. Rather than estimating separate risk models for each variable, we model common risk dynamics through factors and allow heterogeneity to emerge naturally through factor loadings.

Implementation: Three Core Mechanisms

We implement the mean-volatility correlation illustrated above through three interconnected mechanisms in equations 3 and 4:

1. **In-mean effects.** The term $\sum_{k=1}^{K} b_k \tilde{h}_{t-k}$ in equation 3 allows past volatility to directly affect current factor dynamics, following Mumtaz and Theodoridis (2018). Elevated past uncertainty can depress real activity or amplify risk premia.

2. **Endogenous volatility.** Equation 4 allows factor movements to influence volatility evolution through the term $\sum_{j=1}^{Q} d_j F_{t-j}$. Uncertainty responds to economic conditions rather than evolving independently.

3. **Correlated shocks.** We allow non-zero correlation $\Sigma_{\eta e}$ between shocks to factors (e_t) and shocks to their volatilities (η_t). Combined with time-varying volatility (equation 6), shocks that affect factor levels simultaneously influence volatility.

Together, these channels generate the *leverage effect*—the well-documented pattern where negative shocks amplify subsequent volatility in asset prices (Black, 1976; Christie, 1982; Schwert, 1989). Ludvigson et al. (2021) showed that this asymmetry extends to

[1] For this illustration, we shut down the idiosyncratic shocks to focus on factor-driven dynamics.

macroeconomic aggregates: adverse shocks to activity systematically increase uncertainty. In our model, when factors capturing real activity decline, their volatility rises through the mechanisms above, amplifying downside risk and generating the asymmetric tail behavior illustrated in Figure 1. We also include exogenous stochastic volatility in idiosyncratic shocks to capture variable-specific uncertainty, and estimate the factor loading matrix B flexibly to allow heterogeneous risk patterns across variables.

Why Standard DFMs Cannot Capture Risk Dynamics

To appreciate the contribution of our modeling choices, consider how standard dynamic factor models and existing extensions handle uncertainty and risk.

Standard linear DFMs. In a linear factor model with constant volatility, a shock to a factor shifts the mean of the conditional distribution but leaves its width unchanged. The upper and lower tails move symmetrically, and there is no time variation in uncertainty. Such models can characterize average dynamics but cannot capture the risk fluctuations we document.

DFMs with exogenous stochastic volatility. Del Negro and Otrok (2008) and Mumtaz and Surico (2012) introduce time-varying volatility into factor models, allowing the width of conditional distributions to change over time. This captures uncertainty fluctuations. However, the tails still move symmetrically when volatility changes. A volatility shock widens the distribution uniformly, affecting upside and downside risks equally. These models cannot generate *asymmetric* tail risk—the phenomenon where downside risk exhibits greater volatility over time than upside risk, or vice versa.

Comparison to Jurado et al. (2015). The framework in Jurado et al. (2015) is a related approach to measuring uncertainty and its effects. To illustrate the difference, consider a single-factor version of our model:

$$\begin{aligned} F_t &= c_1 + \beta_{11} F_{t-1} + \beta_{12} h_{t-1} + \exp(h_t/2)\epsilon_t \\ h_t &= c_2 + \beta_{21} F_{t-1} + \beta_{22} h_{t-1} + \tau \eta_t \end{aligned} \tag{7}$$

with $\text{Corr}(\epsilon_t, \eta_t) = \zeta$. Jurado et al. (2015) effectively impose $\beta_{12} = \beta_{21} = \zeta = 0$, so that uncertainty evolves independently of factor movements. Time variation in uncertainty is driven primarily by past changes in volatility h_t. In our framework, shocks simultaneously affect both the level and uncertainty of factors ($\beta_{21} \neq 0$, $\zeta \neq 0$), capturing aggregate disturbances with first- and second-moment effects—a channel absent in Jurado et al. (2015). Moreover, our specification allows predictor uncertainty to be endogenous, influenced by

both current shocks and past movements in *both* the factor level F_t and its volatility h_t.[2]

Comparison to Carriero et al. (2018). Carriero et al. (2018) also estimate common macroeconomic and financial uncertainty jointly with mean effects in a factor volatility model. Relative to their work, our framework differs by specifying volatility as originating from the latent factors themselves rather than modeling a separate common volatility factor. This allows us to simultaneously model common fluctuations in levels and volatilities across many observable variables while maintaining, in principle, a clear economic interpretation through named factors, as each factor is associated with a specific observable variable.

2.3 Estimation and Simulation Algorithms

We estimate the model parameters using Gibbs sampling, cycling through conditional posterior distributions for factors, parameters, and volatilities. The key computational challenge is sampling stochastic volatilities given their correlation with factors. We address this using particle Gibbs with ancestor sampling (Andrieu et al., 2010; Lindsten et al., 2014), which maintains good mixing properties even with few particles. The algorithm consists of seven main steps: (1) draw latent factors using the Carter-Kohn algorithm; (2) draw factor loadings and AR coefficients for idiosyncratic shocks; (3) draw idiosyncratic volatilities; (4) draw VAR coefficients; (5) draw variance of volatility shocks via Metropolis step; (6) draw the covariance matrix Σ using the method of Chan and Jeliazkov (2009); (7) draw stochastic volatilities using particle Gibbs. Complete technical details are provided in Appendix B.

We generate 50,000 draws, discard the first 30,000 as burn-in, and retain every 20th draw to minimize serial correlation. Convergence is verified through trace plots and Geweke diagnostics. The resulting 1,000 draws provide reliable posterior inference for all quantities of interest.

A key output of our model is the conditional distribution of each variable at various horizons. These distributions summarize the full range of possible future outcomes, accounting for all sources of uncertainty. For any variable X_i at horizon h, we generate the conditional distribution $p(X_{i,t+h}|X^t)$ by simulating forward from the current state. Each simulation draw proceeds as follows: (1) draw model parameters Θ and current states (F_t, H_t, r_t) from their posterior distribution; (2) simulate factors, volatilities, and idiosyncratic shocks forward h periods using equations 1–4; (3) construct $X_{i,t+h}$ from the

[2] Jurado et al. (2015) include squared factors and factors extracted from squared series as additional predictors, incorporating some stochastic volatility information through this alternative route.

simulated factors and idiosyncratic components. Repeating this process across many posterior draws yields the full conditional distribution, incorporating parameter, state, and innovation uncertainty.

From these distributions, we compute time-varying quantiles to measure downside and upside risk. The volatility of the lower tail (5th percentile) relative to the upper tail (95th percentile) captures the asymmetric risk dynamics we document throughout the paper.

2.4 Data and Model Specification

We use 116 monthly U.S. macroeconomic and financial time series spanning January 1973 to December 2023, covering real activity, inflation, interest rates, asset prices, credit, and disaggregated sectoral indicators. The main data sources are FRED-MD, the Bureau of Economic Analysis's National Income and Product Accounts, Kenneth French's website, and Global Financial Data. Variables are transformed to stationarity and standardized, which is essential for efficient estimation in the factor model framework. A list of variable names is provided in Table A.1.

We estimate seven latent factors (N=7), with six lags of factors and one lag of volatilities in the transition equation (equation 3), and one lag each of volatilities and factors in the volatility equation (equation 4). The number of factors balances model fit with computational tractability; additional factors provide only marginal improvements in explaining variation in the data.[3]

To resolve rotational indeterminacy, we employ the named-factor normalization of Stock and Watson (2016), restricting the first seven rows of the factor loading matrix B to an identity matrix. This pins down each factor through its association with a specific observable variable. We select seven variables that capture distinct dimensions of business cycle and financial fluctuations and exhibit substantial stochastic volatility: Factor 1: Excess bond premium (financial conditions); Factor 2: 1-year Treasury rate (monetary policy); Factor 3: S&P common stock price index (equity markets); Factor 4: Real PCE (aggregate consumption); Factor 5: Real PCE housing and utilities (housing sector); Factor 6: Total non-revolving credit (credit conditions); Factor 7: PCE inflation (price dynamics).

While any seven variables would achieve identification, this particular selection facilitates economic interpretation of the factors and provides strong characterization of the

[3]We consider an extension of the model that allows for Markov regime switching on the intercept terms of the level and volatility equations to account for COVID. This extended model produces similar results for our risk estimates.

common component across our full dataset. Combined with the normalization of the diagonal elements of Σ to equal one, these restrictions fully identify the model.

3 Reading the Tails: Where Does Risk Live?

In this section, we use the estimated model to ask: where does risk live in the economy? We begin by constructing three aggregate risk indices—growth, financial, and inflation—that distill tail risk information across many indicators into economy-wide measures. We then unpack heterogeneity across individual variables and sectors along three dimensions: cross-sectional patterns characterizing the dispersion and clustering of risk asymmetries across variables; the time series evolution of risk for consumption sectors, where quantity and price pairs trace how risk shifts across demand- and supply-driven episodes; and structural determinants of risk for industrial production, where Census Bureau establishment and labor flow data reveal why risk concentrates where it does.

3.1 Aggregate Risk Indices: Growth, Financial, and Inflation

Unlike uncertainty, which is unsigned and reflects overall dispersion, risk is directional: it captures the potential for adverse or favorable outcomes. We construct three separate risk indices—growth, financial, and inflation—aligned with standard classifications of macroeconomic shocks. Adverse **demand shocks** generate downside risk to both growth and inflation. Adverse **supply shocks** produce divergent risks: lower growth, higher inflation. **Financial shocks** generate persistent downside risks to activity through credit spreads and funding conditions.

For each category, we construct two measures capturing downside risk (5th percentile) and upside risk (95th percentile). We use 12-month-ahead conditional distributions for real activity and inflation, reflecting the delayed and persistent transmission of shocks to these variables documented in the VAR literature, and 3-month-ahead distributions for financial variables, which respond immediately to economic stress.

A key advantage of this aggregation approach is robustness. By averaging across many variables, small idiosyncratic shocks cancel out in expectation, isolating the common tail risk signal shared across the economy. At the same time, unusually large idiosyncratic realizations—potentially harbingers of broader sectoral spillovers—remain visible in the index. This signal extraction logic suggests the indices provide a more comprehensive measure of economy-wide tail risk than any individual series, capturing dimensions of risk that standard aggregates may miss.

Each index aggregates the relevant percentile across many standardized indicators using equal weights. **Growth risk**: 60 real activity indicators (labor markets, consumption, housing, sales, inventories, industrial production); Variables are normalized so increases represent improving conditions. **Inflation risk**: 24 price indicators (aggregate and sectoral PCE inflation). **Financial risk**: 32 indicators (equity returns, credit volumes, spreads). Variables are normalized so increases indicate worsening conditions.

Figure 2 plots the indices from 1976 through 2023. Blue lines show the 5th and 95th percentiles; red horizontal lines mark standard normal reference values. All indices are in standard deviation units.

Several patterns emerge. First, sharp spikes during major crises (2008–09 GFC, early 2020 pandemic) reflect synchronized macro-financial risk. Second, there are meaningful asymmetries: growth and financial risk skew toward adverse events; inflation risk tilts upside in the 1970s and post-pandemic, but downside during the GFC. Third, the pandemic period differs from past episodes. Initially, all three indices show simultaneous upside and downside risk (elevated uncertainty). From 2021, risk becomes asymmetric: growth and inflation tilt upside, though persistently so only for inflation. This contrasts with the 1970s (sustained upside inflation risk) and the GFC (downside growth risk). Finally, the 2010s show persistently subdued inflation risk, consistent with below-target dynamics. Empirically, the indices are substantially smoother than variable-specific tail risk measures outside major recessions, while remaining responsive during crises.

3.2 Heterogeneity in Risk Across Variables

The aggregate indices mask substantial variation in the nature and direction of risk across indicators and sectors of the economy. To unpack what drives these aggregate patterns, we characterize each variable's risk profile using two summary statistics computed from its conditional predictive distribution. The first is a measure of tail asymmetry: the log ratio of the standard deviation of the 95th percentile over time to the standard deviation of the 5th percentile over time. Values below zero indicate greater volatility in downside risk relative to upside risk, while values above zero suggest greater volatility in upside risk. Values of -1 (1) mean the lower (upper) tail is twice as volatile. The second is the mean-volatility correlation—the correlation between the conditional mean and conditional volatility of each variable's predictive distribution across time—which captures the leverage effect mechanism generating such asymmetry. Together, these two statistics quantify both the direction of risk and the model mechanism generating it.

Figure 3 plots these two statistics for the 116 sign-normalized variables in our dataset

from July 1976 to June 2019.[4] The aggregate risk indices, shown as pentagons, serve as anchors: the growth risk index sits firmly in the lower-left quadrant, reflecting more volatile downside risk and a strong negative mean-volatility correlation; the inflation and financial risk indices sit in the upper-right quadrant, reflecting more volatile upside risk and a positive mean-volatility correlation, where rising levels are accompanied by rising uncertainty.[5]

The most striking feature of the scatter is that points line up along a positive relationship between the two statistics, cutting across all variable categories. This pattern is informative about the source of tail risk: if idiosyncratic shocks dominated risk dynamics, there would be a cloud of points around the origin with much less range along both axes. The tight relationship over a wide range we observe thus reflects the dominance of common factors in shaping tail risk across the economy, with dispersion around the line capturing heterogeneity in factor exposures across variables.[6]

Within this common structure, there is meaningful heterogeneity across and within categories, reflecting differential factor exposures. Real activity variables cluster in the lower-left quadrant, exhibiting strong negative mean-volatility correlations and pronounced downside risk. Price variables cluster in the upper-right quadrant, exhibiting positive mean-volatility correlations and upside risk. Financial variables are the most dispersed. Interest rate and spread series sit in the upper-right quadrant, while sign-normalized equity variables cluster near the origin or in the lower-left quadrant, as volatility tends to dominate tail risk dynamics. Despite this within-category dispersion, the aggregate financial risk index sits in the upper-right quadrant, as systematic movements in rates and spreads dominate idiosyncratic equity fluctuations in driving the index.

The scatter reveals substantial heterogeneity across variables, both within and across groups. To better understand this variation, we identify the variables with the most pronounced tail asymmetries. Appendix Table A.3 ranks variables with the largest average tail asymmetry in the three categories. On the growth side, the aggregate variables with the most asymmetric downside risk span multiple sectors, including manufacturing, labor markets, and industrial production. Among financial indicators, the strongest asymmetries appear in the VIX and various credit spreads. By contrast, equity returns have

[4]We exclude COVID for this analysis as it was a noneconomic shock that led to unprecedented moves in risk and uncertainty.

[5]For the aggregate risk indices, mean and volatility are constructed by applying the same cross-sectional aggregation used to compute the index quantiles in Section 3.1, averaging the relevant statistics across variables in each category.

[6]In finite samples, realized idiosyncratic shocks can induce spurious mean-volatility correlation; Section 6 confirms that shutting down idiosyncratic components preserves the structure of the scatter, validating the common factor interpretation.

elevated volatility, and this volatility affects both tails.

We now examine sectoral patterns in consumption and production, where heterogeneity plays a central role in understanding risk transmission.

3.3 Risk Dynamics Over History: Consumption Sectors

The cross-sectional patterns in Figure 3 capture average asymmetry over the full sample, masking rich time variation across episodes with very different shock compositions. For consumption sectors, the availability of quantity and price pairs at a disaggregated level allows us to trace how tail risk shifts across demand- and supply-driven regimes.

Figure 4 examines the evolution of tail risk for sectoral personal consumption expenditures (PCE), both quantities and prices. The figure uses heatmaps to plot *excess downside risk* for quantities and *excess upside risk* for prices. Excess downside (upside) risk is computed as the deviation of the 5th (95th) percentile from its historical mean, expressed in standard deviations. Darker shading indicates more pronounced tail risk.

Several findings stand out. First, risk asymmetries are highly episodic and sector-specific. While crisis periods tend to see a generalized increase in risk across sectors, the magnitude and persistence vary considerably. Second, the 1970s were characterized by more acute, persistent, and broad-based upside risk to inflation compared to the post-pandemic period. During the 1970s, nearly all consumption categories exhibited elevated and sustained inflation-at-risk. In contrast, during late 2020 and early 2021, inflation risk spiked sharply but remained elevated only in a subset of sectors, particularly those related to goods consumption such as motor vehicles and furnishings.

Third, sector-specific risks are clearly visible. For instance, housing-related consumption shows pronounced downside quantity risk and upside price risk during the mid-2000s housing boom and subsequent bust. Energy consumption presents an interesting case: in the 1970s, price controls muted upside price risk relative to other sectors, whereas during the COVID period, energy prices exhibited sharp asymmetries reflecting the full impact of the demand collapse and subsequent recovery alongside supply constraints.

Fourth, the dynamics of risk in consumption quantities mirror, in opposite direction, those in prices during certain episodes. During supply-constrained periods such as the 1970s or specific phases of the pandemic recovery, quantity downside risk coincides with price upside risk. By contrast, during the GFC, quantities exhibit pronounced downside risk while prices show muted or even downside inflation risk, consistent with demand-driven dynamics.

3.4 Industry Characteristics and Risk

Finally, we ask: *Why does risk live where it does?* We focus on industrial production, where Census Bureau microdata on establishment and labor market dynamics at the three-digit NAICS level allow us to link tail risk patterns to observable structural characteristics of industries. Figure 5 presents three panels that answer this question in two complementary ways: through external validation using industry dynamics Census Bureau data not used in estimation, and through the model mechanism linking tail asymmetry to differential exposure to financial conditions.

We leverage BDS microdata, which provides annual data on job flows and establishment dynamics at the three-digit NAICS level. We focus on two measures of industry dynamism: the net job creation rate (job creation minus destruction) and the net establishment entry rate (entry minus exit), calculated annually from 1978 to 2019.[7] For both variables, we compute the Kelley skewness by industry over this period.

The first two panels document that industries with more downside output risk also exhibit more left-skewed labor market and firm dynamics. Sectors such as nonmetallic mineral products and chemicals display pronounced left skewness in both net job creation and net establishment entry rates—they are more likely to shed jobs and close establishments during downturns than to expand during recoveries—and these same industries exhibit the most volatile downside risk in production. Since BDS data are not used in estimating the model, this correlation provides external validation of the risk patterns the model generates, and points to labor market adjustment and business formation characteristics as structural sources of asymmetry.

The third panel addresses the model mechanism. Industries with more negative loadings on the financial conditions factors—factors 1 and 3, the only factors whose loadings are individually and jointly statistically significant predictors of tail asymmetry across IP sectors—exhibit more pronounced downside risk. This association points to differential exposure to financial conditions as a key driver of heterogeneity in tail risk across industries. Nonmetallic mineral products anchors the bottom of the distribution, with the most negative financial factor loadings and the most volatile downside risk, while mining sits near the top with near-zero or positive loadings and relatively symmetric risk profiles. These results underscore that asymmetries in risk vary systematically across sectors

[7]The job destruction (creation) rate is calculated as the count of all employment losses (gains) from contracting (expanding) and closing (opening) establishments, divided by the average of employment for times t and $t-1$. The establishment exit (entry) rate is defined as the count of establishments exiting (born) during the year, divided by the average of establishments for times t and $t-1$. The results remain robust if we include data through the end of our sample but we exclude 2020 because of COVID.

and industries, and reflect deeper structural characteristics of production and industry dynamics.

4 Factor Estimates and Model Mechanisms

Having documented the behavior of risk across variables and over time, we now examine the model mechanisms that generate these patterns. The analysis proceeds in three parts. First, we examine the estimated paths of the latent factors and their volatilities over time. Second, we explore the joint dynamics of means and volatilities, focusing on the within-factor correlation that gives rise to asymmetric risk (the leverage effect) and the cross-factor correlation that shapes co-movement. Third, we examine how factors load onto variables, generating the heterogeneity in sectoral risk documented in Section 3.

4.1 Latent Factors and Volatility Over Time

The left column of Figure 6 plots the median estimates of four of the seven latent factors F_t alongside the specific observable variable used in the normalization described in Section 2. The shaded regions indicate 90 percent posterior credible sets. The right column displays the estimated log volatilities of the same factors. The remaining factors and volatilities are reported in Appendix Figure A.7.

The factors are estimated precisely and admit clear economic interpretations. The first factor captures financial conditions, closely tracking the excess bond premium. The fourth factor reflects aggregate consumption dynamics, mirroring real personal consumption expenditures. The sixth factor captures credit market conditions. The seventh factor is associated with PCE inflation dynamics. While these normalizations aid interpretation, each factor reflects information common across a broader set of variables beyond its anchoring observable. The dynamic factor model distinguishes between movements that are broadly shared, the common component captured by the factors, and those that are specific to individual series, the idiosyncratic component visualized as the gap between the factors and the observable series plotted in each panel.

Volatility varies markedly over time and is strongly countercyclical, rising during recessions and subsiding in expansions. Volatility spikes are particularly pronounced during well-known episodes of economic and financial stress: the early 1980s recessions, the GFC, and the onset of the COVID-19 pandemic. While some movements are factor-specific, there is substantial co-movement across volatilities, indicating systemic shifts in aggregate uncertainty.

4.2 Mean–Volatility Interactions and the Leverage Effect

As discussed in Section 2, a central mechanism generating asymmetric risk in the model is the leverage effect, the interaction between the level of each factor and its volatility. In our setting, this manifests as correlation between factor movements and changes in their conditional variance. Figure 7 sheds light on this mechanism by visualizing the full correlation structure between the means and volatilities of the seven latent factors, computed from their smoothed values over the full sample. This heatmap reveals several economically meaningful patterns.

First, factors associated with financial conditions (factors 1 and 2) and inflation (factor 7) display strong positive mean–volatility correlations. When financial conditions tighten or inflation rises, uncertainty increases. This positive correlation generates more volatile upper tails for variables that load positively on these factors—consistent with the inflation-at-risk and spread-at-risk patterns documented in Section 3. Second, the factor related to credit (factor 6) exhibits negative mean–volatility correlations, consistent with a leverage effect. Negative shocks to credit raise volatility, amplifying downside risk. This mechanism, along with the fact that many real activity variables load negatively on the financial and inflation factors, explains the prevalence of growth-at-risk behavior—more volatile lower tails—observed across these indicators in Section 3.

Third, the heatmap reveals widespread co-movement in volatility across factors, indicated by the predominantly positive correlations in the volatility–volatility block (green shading). This co-movement of uncertainty drives the common risk patterns observed across many variables during crisis episodes. Finally, the means of certain factors, particularly the financial factor (factor 1) and the credit factor (factor 6), are strongly correlated with the volatilities of other factors. This highlights their central roles in transmitting and amplifying macroeconomic risk across the system.

Appendix Figure A.8 provides additional granularity, showing that many of these mean volatility correlations are precisely estimated and vary across subsamples. For example, factors related to interest rates (factor 2) and inflation (factor 7) exhibit markedly more negative mean–volatility correlations during the GFC than in the pre-1983 period. This shift reflects the transition from an era dominated by upside inflation risk to one characterized by financial stress and disinflationary pressures. These time-varying correlations underscore the state-dependent nature of risk transmission in the model.

4.3 Factor Loadings and the Origins of Sectoral Risk

The heterogeneity in risk across sectors documented in Section 3 has a structural explanation: it arises from the interaction between heterogeneous factor loadings and the mean-volatility correlations documented in Section 4.2. The factor loading matrix B determines how shocks to factors, and their associated volatility dynamics, transmit to individual observables, shaping not just the level of risk exposure but its direction—whether a variable faces predominantly downside or upside risk—and its magnitude across sectors.

Figure 8 displays the median posterior estimates of factor loadings for personal consumption expenditure (PCE) sectors. The top panel shows results for PCE quantities; the bottom panel for PCE prices. Colored bars highlight the four most influential factors on average in each case, determined by the magnitude of loadings and their contribution to cross-sectional variation.[8]

Several economically interpretable patterns emerge. For quantities, virtually all sectors load positively on factor 4, which captures aggregate real PCE dynamics. However, the magnitude varies substantially—explaining differences in volatility and cyclicality across consumption categories. Critically, some sectors also load negatively on the financial conditions factor (factor 1). For example, durable goods consumption has a large negative loading, implying that tightening financial conditions translate directly into elevated downside risk for this sector. This mechanism explains the pronounced growth-at-risk behavior of durables documented in Section 3. Similarly, certain sectors load negatively on the inflation factor (factor 7), meaning that inflationary pressures depress real consumption in these categories.

For PCE prices, most sectors load positively on the inflation factor (factor 7), as expected. However, the magnitude varies considerably across categories. Many sectors also load positively on the financial factor (factor 1), suggesting that tighter financial conditions are associated with increased upside price risk—likely reflecting supply-side disruptions or demand inelasticity in certain goods. The combination of positive loadings on factors with positive mean-volatility correlations (factors 1 and 7) generates the inflation-at-risk patterns observed in Section 3.

These loading patterns provide the foundation for the asymmetries documented in Section 3. When combined with the mean-volatility correlations shown in Figure 7, they explain why certain sectors exhibit pronounced downside risk, others exhibit upside risk, and still others display relatively symmetric distributions. The model's ability to gener-

[8]Factor loadings for industrial production (IP) by industry are reported in Appendix Figure A.9. See Section 3 for a discussion of the relationship between these loadings and tail asymmetry in IP.

ate this rich heterogeneity parsimoniously, through the interaction of seven factors with heterogeneous loadings, is a key advantage over approaches that model risk separately for each variable. In particular, the negative loadings of real activity variables—both consumption and IP—on the financial conditions factor (factor 1), combined with that factor's positive mean-volatility correlation, provide a unified account of why financial conditions predict left-tail growth risk as documented by Adrian et al. (2019), and show that this pattern extends pervasively across real activity variables.

5 Risk Propagation: Validating the Mechanisms

Sections 3 and 4 documented substantial heterogeneity in tail risk across sectors and showed that this heterogeneity correlates with factor loadings and mean-volatility correlations. We now verify that these statistical relationships reflect genuine economic mechanisms by tracing how a financial shock propagates through the model. We do the same for an inflation shock in Appendix C.

Beyond validation, this exercise demonstrates a practical application of the framework: *granular risk analysis through counterfactual scenarios*. By conditioning on different economic states and varying shock magnitudes, the model can assess how tail risks evolve across sectors under alternative stress scenarios—a tool relevant for both policy analysis and risk management. We begin with the simplest such exercise: a financial shock during the GFC. Then, we examine counterfactual sectoral risk scenarios if the realized financial shocks had not occurred in the GFC and the realized inflation shocks had not occurred in the Great Inflation.

5.1 Factor and Volatility Responses

Figure 9 traces the propagation of a one standard deviation shock to financial conditions (factor 1) in October 2008. This shock is generated assuming a Cholesky decomposition and is ordered first. The shock generates rich co-movement across factors and volatilities, consistent with the systemic nature of financial stress. Financial conditions tighten persistently for approximately two years, capturing the prolonged credit market dysfunction observed during the crisis. Real activity (factor 4) declines sharply on impact and recovers only gradually, mirroring the sluggish pace of the post-crisis recovery. Credit conditions (factor 6) exhibit a hump-shaped decline, peaking around six months after the initial shock as financial stress transmits to lending standards and credit availability. The inflation response (factor 7) is particularly revealing: an initial rise—likely reflecting

cost-push pressures or temporary supply disruptions—followed by a persistent decline as weak demand dominates. This pattern aligns closely with observed inflation dynamics during the crisis, when commodity price spikes gave way to deflationary concerns.

Critically, log volatility increases substantially for financial and credit factors, amplifying uncertainty precisely when economic conditions deteriorate. This dual movement generates pronounced downside risk—the growth-at-risk patterns observed in Section 3. The inflation factor's volatility shows little response, consistent with its more muted mean-volatility correlation during this period.

5.2 Sectoral Heterogeneity: Plausible or Surprising?

Figure 10 shows how the same financial shock generates strikingly different distributional responses across sectors. We present results for six sectors chosen to span the range of cyclicality and financial sensitivity in the economy. The heterogeneity is both economically plausible and quantitatively large.

Industrial production. Nonmetallic mineral products—which includes cement, concrete, and other construction materials—exhibits extreme downside risk sensitivity. The distribution shifts dramatically leftward with the 5th percentile plummeting while the 95th percentile barely moves. This sector's production is closely tied to construction activity, which depends heavily on credit availability. When financial conditions tighten, construction projects are postponed or cancelled, causing severe contractions in demand for these materials. The sector's large negative loading on the financial factor captures this transmission channel. In contrast, electric power generation, while also showing downside risk, exhibits a much more muted response. As a regulated utility with stable demand and limited direct credit dependence, electricity production is partially insulated from financial shocks.

Thus, both sectors experience downside risk in response to tighter financial conditions—consistent with the general growth-at-risk pattern—but the *magnitude* differs dramatically, validating the model's ability to capture heterogeneity in cyclical sensitivity across industrial sectors.

Personal consumption: quantities. Furnishings and durable household equipment, big-ticket items financed through credit, exhibit sharp downside quantity risk. The 5th percentile drops dramatically, consistent with consumers postponing large purchases when credit tightens or economic uncertainty rises. Food and beverages, by contrast, show minimal response—unsurprising for necessity goods with low income and credit elasticity. The model characterizes consumption heterogeneity that accords with basic economic

intuition about durables versus nondurables.

Personal consumption: prices. The price responses reveal additional nuance. Furnishings prices shift slightly to the *right* with increased dispersion despite falling quantities, reflecting inflation-at-risk. This pattern is consistent with supply chain disruptions or cost-push pressures dominating demand effects for import-intensive durables during the crisis. Food prices shift left as weak demand dominates. These divergent price dynamics across sectors—inflationary for some goods and deflationary for others—align with the sectoral inflation heterogeneity observed during 2008-2009.

Taking Stock. The results pass a basic plausibility test: cyclical, credit-sensitive sectors exhibit large asymmetric responses; acyclical necessities remain stable; price and quantity dynamics accord with sectoral characteristics. More importantly, the magnitudes are economically meaningful. They are not just statistically significant but also quantitatively large enough to matter for sectoral composition effects on aggregate risk. An economy more heavily weighted toward construction materials and durables would face substantially greater aggregate tail risk than one weighted toward utilities and food, a composition effect invisible in aggregate uncertainty measures but revealed here.

5.3 Time-Varying Risk Dynamics: Supply versus Demand

The risk patterns documented in Section 3 vary over time in economically interpretable ways. Figure 11 illustrates this for sectoral inflation risk using two contrasting historical periods: the Great Inflation (1978-82) and the GFC (2007-11). For each period, we construct counterfactuals by shutting down realized shocks to the financial factor and the inflation factor separately over these two episodes. We assume a Cholesky decomposition with the shock to the financial factor ordered first and inflation factor ordered second. We emphasize that the goal of this exercise is not structural identification—which would require a dedicated framework beyond the scope of this paper—but rather to illustrate how the model's factor structure generates distinct risk patterns under different shock configurations. Appendix D provides additional details about the implementation.

The contrast is striking. During the Great Inflation (left panel), sectoral price inflation exhibits pronounced upside risk—the upper tail is substantially more volatile than the lower tail across most categories. Shutting down inflation shocks (open triangles) eliminates this asymmetry, while shutting down financial shocks (open circles) has minimal effect. During the GFC (right panel), the pattern reverses: prices show downside risk alongside quantities, characteristic of demand-driven dynamics. Here, financial shocks (open circles) account for most tail asymmetries, while inflation shocks play a minor role.

This regime-dependent behavior demonstrates that the framework captures economically distinct shock types generating different risk patterns, which is the kind of state-dependence predicted by models with occasionally binding constraints or regime-switching dynamics. The same model structure produces both inflation-at-risk and growth-at-risk when supply shocks dominate, while generating primarily growth-at-risk when demand shocks dominate.

6 Model Fit and Robustness

Having documented the pervasive heterogeneity in tail risk across variables and over time, we now assess how well the model fits the data and captures the dynamics of conditional distributions. We present three complementary pieces of evidence. First, we examine the share of variance in observable variables explained by the common factors. Second, we show that the heterogeneity and structure of risk that our model generates are driven by the loading structure on the common factors. Finally, we compare the model's implied conditional distributions to those from quantile regressions, a popular alternative in the tail risk literature, both in-sample and out-of-sample.

6.1 Variance Decomposition: How Much Do Factors Explain?

The model's ability to generate time-varying and asymmetric risk across a large set of variables depends critically on the relevance of the common factors. Panel A of Table 1 reports the share of variance in selected aggregate series explained by the common component—calculated as the ratio of variance attributable to factors relative to total variance in the data. Panel B extends this analysis to sectoral series.

At the aggregate level, the factors account for substantial variation in many real activity and price indicators. Industrial production, real PCE, and PCE inflation—especially in services—are well explained by the common component, with R^2 values ranging from 0.5 to 0.9. For example, total industrial production has an R^2 of 0.90, indicating that 90 percent of its variation is driven by common dynamics. Real PCE displays an R^2 of 0.62, while services inflation shows an R^2 of 0.74. In contrast, some credit aggregates and specific financial variables display more idiosyncratic behavior with lower R^2 values. These results are consistent with earlier findings that dynamic factor models perform best for real activity and inflation (e.g., Stock and Watson, 2016).

At the sectoral level, explanatory power varies widely but remains economically meaningful. The common component explains, on average, 34 percent of fluctuations in in-

dustrial production across sectors, but this share ranges from under 10 percent to over 70 percent across industries. Sectors such as motor vehicles and plastics and rubber products, which are closely tied to aggregate demand and financial conditions, exhibit higher R^2 values. For sectoral real PCE, the average is 20 percent, with furnishings and durable household equipment showing particularly strong common component influence ($R^2 = 0.44$). PCE inflation by category displays an average R^2 of 19 percent, with services inflation more strongly driven by common factors than goods inflation. Finally, equity returns are heavily influenced by the factors, with an average explanatory share of 70 percent across industry portfolios, consistent with the aggregate nature of shocks priced in financial markets.

These results highlight that the model successfully captures common variation across a broad set of macroeconomic and financial indicators. This structure, generated by substantial commonality coupled with meaningful idiosyncratic variation, is central to the model's ability to generate the heterogeneous yet systematically correlated risk patterns documented in Section 3.

6.2 Generating Heterogeneity and Structure in Risk

The model generates the heterogeneity and structure in risk documented in Section 3 through the loadings on the common factors. Figure 12 drills down on this relationship. The top panel reproduces Figure 3, which shows the relationship between mean-volatility correlation and tail risk across all variables. The second panel shuts off the idiosyncratic components of the variables. Comparing the first and second panels, the qualitative patterns continue to hold. Growth variables congregate in the bottom-left quadrant, inflation variables in the upper-right, and financial variables spread across both. The common component therefore plays a critical role in generating the risk patterns we see. The bottom panel randomly reshuffles the factor loadings for each variable—reassigning each variable's loadings across factors in a single random draw—while keeping the idiosyncratic components off. Now, the variables in growth, inflation, and financial spread across the two quadrants. The behavior of recreational goods and vehicles and transportation services inflation is instructive. Both inflation series are in the upper-right quadrant in the first two panels, showing the usual inflation-at-risk behavior. Upon reshuffling of the loadings, the two series switch to the bottom-left quadrant. This exercise illustrates that the factor loading structure, combined with the factors themselves, generates the structure of risk in the model.

6.3 Comparison to Quantile Regressions

To benchmark the performance of our model, we compare its conditional distributions to those generated by quantile regressions, a widely used alternative for modeling tail risk. Following Adrian et al. (2019), we estimate quantile regressions of the form:

$$x_{t+h}^{(h)} = \beta_{0,\tau} + \beta_{1,\tau} x_t + \beta_{2,\tau} NFCI_t + \epsilon_{\tau,t+h}, \tag{8}$$

where $x_{t+h}^{(h)}$ denotes the variable of interest at horizon h, and $\tau \in \{5\%, 25\%, 75\%, 95\%\}$. The National Financial Conditions Index (NFCI) is included to capture time-varying financial conditions. We fit the Azzalini and Capitanio (2003) skewed-t distribution to the resulting quantile forecasts to generate full conditional densities.

We examine in-sample forecasting performance and conduct a recursive out-of-sample forecast evaluation using an expanding window approach. The in-sample forecasts run monthly from July 1976 through the end of 2023. For the out-of-sample exercise, we begin with forecasts conditional on data through December 2006, then expand the sample by three months and reestimate both models to generate the next forecast through the end of 2019. To evaluate tail performance specifically, Table 2 reports ratios of quantile-weighted continuous ranked probability scores (qwCRPS) (Gneiting and Ranjan, 2011), a scoring rule that places greater weight on the accuracy of tail predictions relative to the center of the distribution, following Carriero et al. (2024).[9] We use the quantile regression forecasts as the baseline, so ratios below 1 indicate superior tail performance by the dynamic factor model. The first column shows results for the full evaluation sample; the second and third columns break down results into recession and non-recession periods.

The dynamic factor model forecasts real activity and inflation variables more accurately at the tails, while the quantile regression approach performs better for financial indicators. The latter result is not surprising given that the quantile regression uses the NFCI, a financial factor that aggregates over a hundred series, as a predictor variable. These rankings are broadly consistent across our in-sample and out-of-sample exercises. The factor model performs relatively better at the 3-month horizon but remains competitive at the 12-month horizon. In recessions, the dynamic factor model continues to perform well and in some cases—such as 3-month ahead IP growth and right-tail PCE growth forecasts in-sample—even outperforms the quantile regressions.

These results suggest that the two approaches have complementary strengths. The

[9]The qwCRPS is defined as $qwCRPS_t = \frac{2}{J-1} \sum_{j=1}^{J-1} v(\tau_j) QS_{\tau_j,t}$ where $QS_{\tau_j,t}$ is the quantile score of Giacomini and Komunjer (2005). The left qwCRPS uses the weight function $v(\tau_j) = (1-\tau_j)^2$ and the right qwCRPS uses $v(\tau_j) = \tau_j^2$ where $\tau_j = 0.05, 0.1, ..., 0.95$ quantiles and $J = 20$.

dynamic factor model excels at capturing tail risk dynamics for variables with strong common components—real activity and inflation. Quantile regressions, which directly condition on financial conditions through the NFCI, perform better for variables that respond immediately to financial stress. Importantly, both approaches produce qualitatively similar conditional distributions for key variables, validating the main risk patterns documented in this paper.

7 Conclusions

This paper shows that asymmetric tail risk is a systematic, pervasive feature of macroeconomic dynamics. Across 116 U.S. time series spanning five decades, tail risk varies dramatically: real activity faces predominantly downside risk, prices exhibit regime-dependent asymmetries, and financial variables show upside risk during stress. Rather than requiring separate explanations, a unified framework with endogenous stochastic volatility in common factors generates growth-at-risk, inflation-at-risk, and sectoral heterogeneity simultaneously. Heterogeneity reflects differential factor exposures, not idiosyncratic noise.

The framework offers several directions for future work. First, it is compatible with structural identification strategies that could enable formal policy counterfactuals—assessing how alternative monetary or fiscal policies reshape tail risk across sectors. Second, extending to other economies could reveal whether these mechanisms operate similarly internationally or whether country-specific institutions alter propagation. Third, incorporating time-varying parameters beyond volatility could address whether factor relationships remain stable across long historical periods. Finally, the distributional measures we provide offer targets for calibrating and testing dynamic models with nonlinear features: models with occasionally binding constraints, financial frictions, or regime-switching dynamics that predict asymmetric tail behavior but currently lack comprehensive empirical benchmarks.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289.

Amburgey, A. and M. W. McCracken (2025): "Growth-at-Risk is Investment-at-Risk," Working Papers 2023-020, Federal Reserve Bank of St. Louis.

Andrieu, C., A. Doucet, and R. Holenstein (2010): "Particle Markov chain Monte Carlo methods," *Journal of the Royal Statistical Society Series B*, 72, 269–342.

Azzalini, A. and A. Capitanio (2003): "Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t-Distribution," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65, 367–389.

Banbura, M., D. Giannone, and L. Reichlin (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25, 71–92.

Bernanke, B., J. Boivin, and P. S. Eliasz (2005): "Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach," *The Quarterly Journal of Economics*, 120, 387–422.

Black, F. (1976): "Studies of Stock Market Volatility Changes," *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Caldara, D., C. Scotti, and M. Zhong (2021): "Macroeconomic and Financial Risks: A Tale of Mean and Volatility," International Finance Discussion Papers 1326, Board of Governors of the Federal Reserve System (U.S.).

Canova, F. and M. Ciccarelli (2013): "Panel vector autoregressive models: A survey," *VAR models in macroeconomics–New developments and applications: Essays in honor of Christopher A. Sims*, 32, 205–246.

Canova, F. and M. O. Ravn (2000): "The macroeconomic effects of German unification: Real adjustments and the welfare state," *Review of Economic Dynamics*, 3, 423–460.

Canova, F. and L. Sala (2009): "Back to square one: Identification issues in DSGE models," *Journal of Monetary Economics*, 56, 431–449.

Carriero, A., T. Clark, and M. Marcellino (2018): "Measuring Uncertainty and Its Impact on the Economy," *Review of Economics and Statistics*, 100, 799–815.

Carriero, A., T. E. Clark, and M. Marcellino (2024): "Capturing Macro-Economic Tail Risks with Bayesian Vector Autoregressions," *Journal of Money, Credit and Banking*, 56, 1099–1127.

Carter, C. K. and R. Kohn (1994): "On Gibbs sampling for state space models," *Biometrika*, 81, 541–53).

Chan, J. C.-C. and I. Jeliazkov (2009): "MCMC Estimation of Restricted Covariance Matrices," *Journal of Computational and Graphical Statistics*, 18, 457–480.

Christie, A. A. (1982): "The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects," *Journal of Financial Economics*, 10, 407–432.

Davidson, S. N., C. Hou, and G. Koop (2025): "Investigating Economic Uncertainty Using Stochastic Volatility in Mean VARs: The Importance of Model Size, Order-Invariance and Classification," *Journal of Business & Economic Statistics*, 43, 992–1007.

Del Negro, M. and C. Otrok (2008): "Dynamic factor models with time-varying parameters: measuring changes in international business cycles," Tech. rep.

Favara, G., S. Gilchrist, K. F. Lewis, and E. Zakrajšek (2016): "Updating the Recession Risk and the Excess Bond Premium," FEDS Notes 2016-10-06, Board of Governors of the Federal Reserve System (U.S.).

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, and M. Uribe (2011): "Risk matters: The real effects of volatility shocks," *American Economic Review*, 101, 2530–2561.

Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007): "Estimating macroeconomic models: A likelihood approach," *The Review of Economic Studies*, 74, 1059–1087.

Giacomini, R. and I. Komunjer (2005): "Evaluation and Combination of Conditional Quantile Forecasts," *Journal of Business & Economic Statistics*, 23, 416–431.

Gneiting, T. and R. Ranjan (2011): "Comparing Density Forecasts Using Threshold- and Quantile-Weighted Scoring Rules," *Journal of Business & Economic Statistics*, 29, 411–422.

Gorodnichenko, Y. and S. Ng (2017): "Level and Volatility Factors in Macroeconomic Data," *Journal of Monetary Economics*, 91, 52 – 68.

Guerrieri, V. and G. Lorenzoni (2017): "Credit crises, precautionary savings, and the liquidity trap," *The Quarterly Journal of Economics*, 132, 1427–1467.

Jacquier, E., N. G. Polson, and P. E. Rossi (2002): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business & Economic Statistics*, 20, 69–87.

Jurado, K., S. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Kiley, M. T. (2022): "Unemployment risk," *Journal of Money, Credit and Banking*, 54, 1407–1424.

Lindsten, F., M. I. Jordan, and T. B. Schön (2014): "Particle Gibbs with Ancestor Sampling," *Journal of Machine Learning Research*, 15, 2145–2184.

Ludvigson, S. C., S. Ma, and S. Ng (2021): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Journal: Macroeconomics*, 13, 369–410.

López-Salido, D. and F. Loria (2024): "Inflation at risk," *Journal of Monetary Economics*, 145, 103570, inflation: Expectations Dynamics October 14-15, 2022.

Mumtaz, H. and P. Surico (2012): "Evolving International Inflation Dynamics: World and Country-Specific Factors," *Journal of the European Economic Association*, 10, 716–734.

Mumtaz, H. and K. Theodoridis (2018): "The Changing Transmission of Uncertainty Shocks in the U.S." *Journal of Business & Economic Statistics*, 36, 239–252.

Opschoor, D., D. v. Dijk, and P. H. Franses (2025): "Heterogeneity in Manufacturing Growth Risk," *Journal of Money, Credit and Banking*, n/a.

Schwert, G. W. (1989): "Why Does Stock Market Volatility Change Over Time?" *The Journal of Finance*, 44, 1115–1153.

Stock, J. and M. Watson (2016): "Chapter 8 - Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics," Elsevier, vol. 2 of *Handbook of Macroeconomics*, 415–525.

Stock, J. H. and M. W. Watson (2002): "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business & Economic Statistics*, 20, 147–62.

——— (2025): "Recovering from COVID," *Brookings Papers on Economic Activity*, 56, 297–374.

Table 1: Volatility Decomposition: Common and Idiosyncratic Components

		Standard Deviation (Percent)			
Variable	Category	Data	Common	Idio	R^2
Panel A: Aggregate Series					
IP	Total	0.97	0.92	0.28	0.90
Real PCE	Total	0.86	0.67	0.32	0.62
PCE Inflation	Total	0.25	0.17	0.18	0.48
	Services	0.21	0.18	0.09	0.74
Excess bond premium	-	0.55	0.25	0.43	0.21
Equity Prices	S&P Common	3.61	2.13	2.69	0.35
Credit	Total Non-Rev.	0.68	0.21	0.65	0.10
Panel B: Sectoral Series					
IP	Average	2.17	1.20	1.70	0.34
	Median	1.77	0.96	1.30	0.38
	Minimum	0.95	0.47	0.65	0.08
	Maximum	8.49	5.57	5.49	0.72
	Standard Dev	1.68	1.03	1.16	0.17
Real PCE	Average	2.07	0.89	1.65	0.20
	Median	2.03	0.50	1.65	0.15
	Minimum	0.72	0.13	0.63	0.01
	Maximum	5.74	3.59	4.27	0.44
	Standard Dev	1.23	0.85	0.89	0.14
PCE Inflation	Average	0.72	0.25	0.67	0.19
	Median	0.47	0.18	0.44	0.17
	Minimum	0.24	0.06	0.19	0.05
	Maximum	4.29	0.95	4.06	0.33
	Standard Dev	0.98	0.21	0.93	0.09
Equity Prices	Average	5.28	4.37	2.76	0.70
	Median	5.09	4.71	2.59	0.76
	Minimum	3.94	2.52	1.22	0.38
	Maximum	6.56	5.47	5.26	0.92
	Standard Dev	0.79	0.91	1.15	0.20

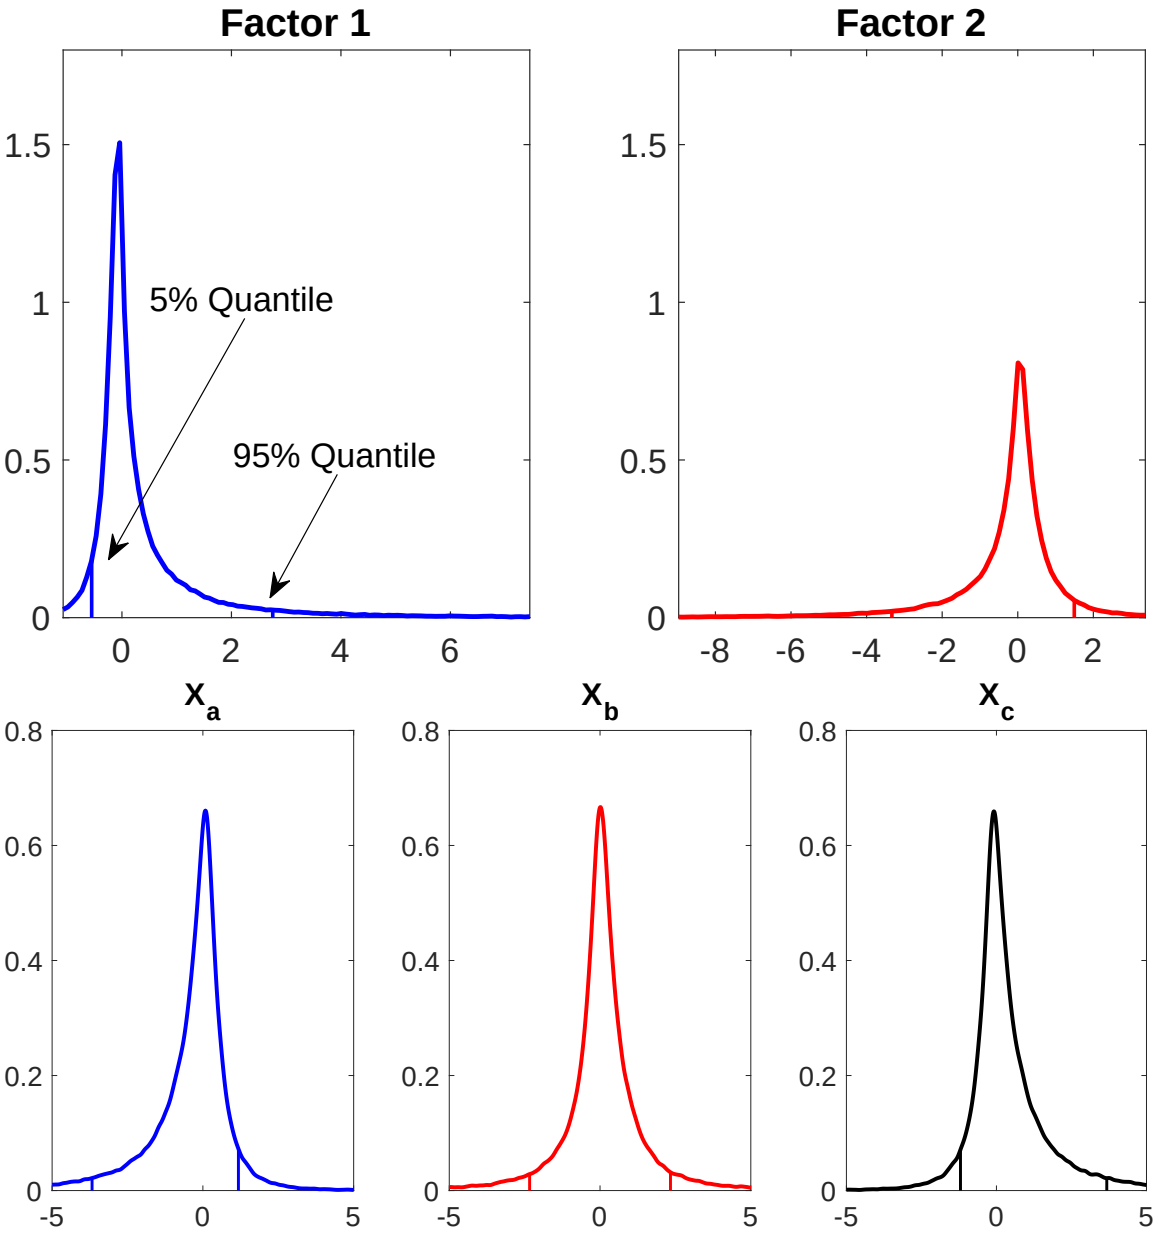
Note: Standard deviations are in percent. For each series, this table shows the standard deviation of (i) the data, (ii) implied by conditioning only on the common component spanned by the factors, (iii) implied by only the idiosyncratic component, and the R^2 of the common component, calculated as the variance of the common component divided by variance of the data. Panel A reports selected aggregate series; Panel B reports summary statistics across sectoral series.

Table 2: Quantile-Weighted CRPS Ratios: Factor Model vs. Quantile Regressions

		In-Sample			Out-of-Sample		
		Full Sample (1976–2023)	Recession	Ex-Recession	Full Sample (2007–2019)	Crisis (2007–2009)	Post-Crisis (2009–2019)
Panel A: 3 Months Ahead							
IP Growth	Left	0.84	0.80	0.85	0.89	0.91	0.87
	Right	0.77	0.69	0.80	0.87	0.83	0.89
Real PCE Growth	Left	0.92	0.94	0.91	0.98	0.93	1.02
	Right	0.82	0.68	0.86	0.90	0.93	0.88
Inflation	Left	0.83	0.83	0.83	0.88	0.83	0.94
	Right	0.82	0.86	0.81	0.94	0.93	0.96
EBP	Left	1.07	1.08	1.07	1.01	0.99	1.03
	Right	1.06	1.06	1.06	1.14	1.21	1.07
Panel B: 12 Months Ahead							
IP Growth	Left	0.94	0.98	0.93	0.93	1.03	0.80
	Right	0.93	0.82	0.96	0.97	1.07	0.89
Real PCE Growth	Left	1.00	1.02	1.00	1.01	1.09	0.91
	Right	0.89	0.72	0.95	0.97	1.03	0.92
Inflation	Left	0.84	0.85	0.84	0.76	0.64	0.92
	Right	0.83	0.74	0.85	0.77	0.61	0.89
EBP	Left	1.07	1.03	1.08	0.98	0.89	1.15
	Right	1.05	0.98	1.07	1.03	1.02	1.06

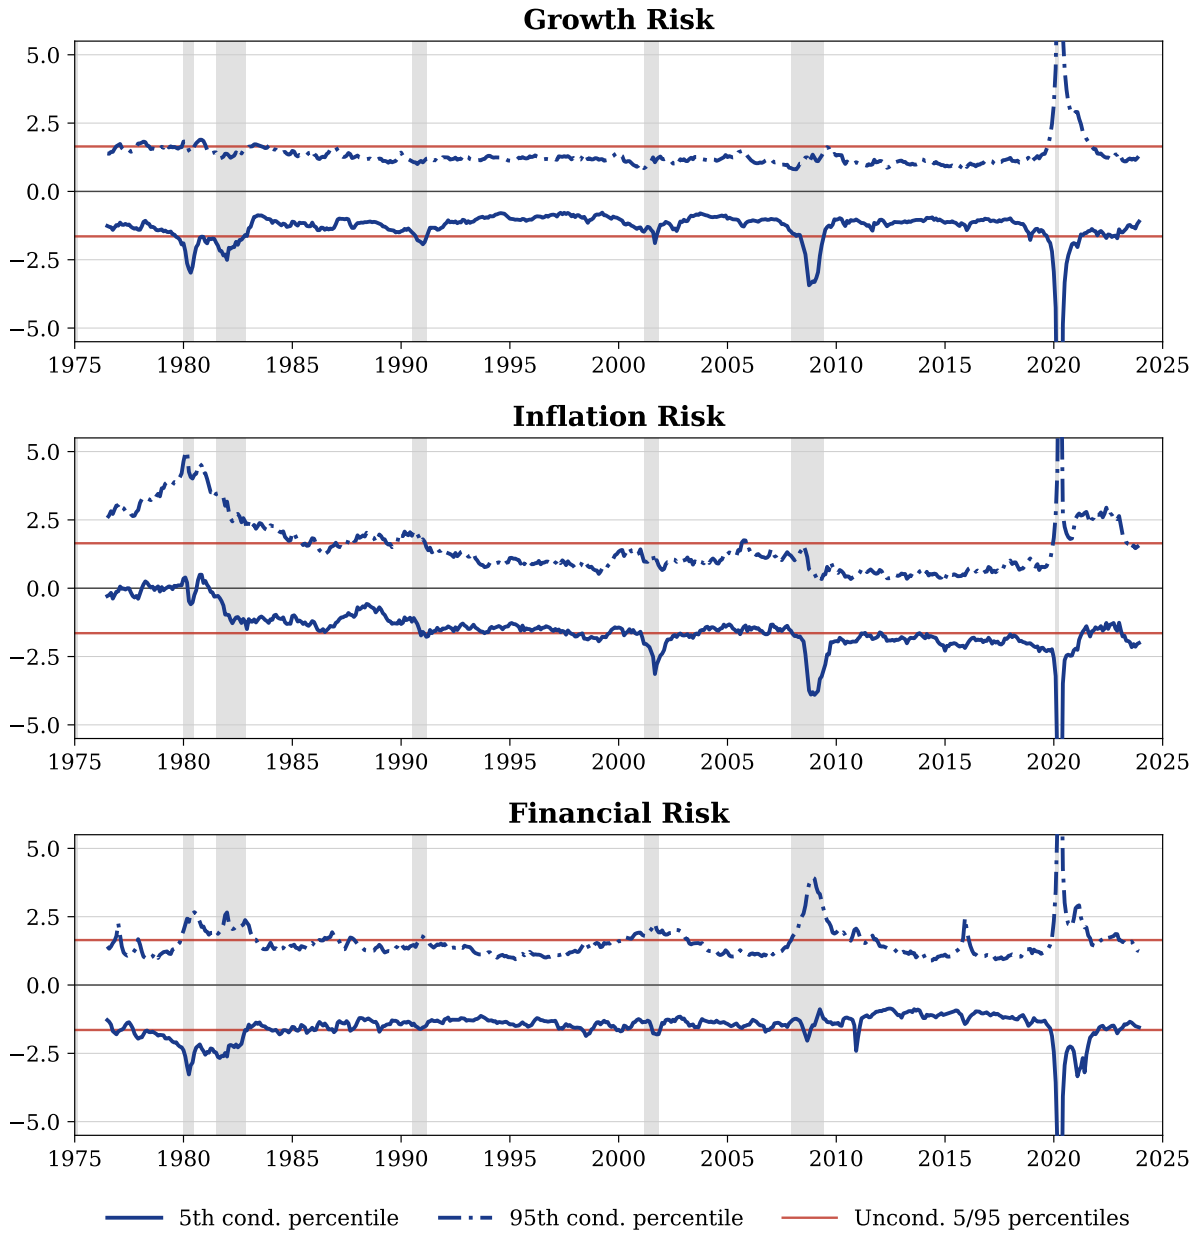
Note: The table reports the ratio of the dynamic factor model's quantile-weighted continuous ranked probability scores (CRPS) to that of the quantile regression baseline (Factor/QR). Values below 1 indicate the factor model forecasts better (lower loss). Left and Right refer to left-tail and right-tail weighted CRPS, emphasizing downside and upside forecast accuracy, respectively. For the in-sample results, full sample spans July 1976 through September 2023 (3 months ahead) and December 2022 (12 months ahead). Recession and Ex-Recession split the sample by NBER recession months. Out-of-sample forecasts use a recursive expanding window with the first forecasts made using data estimated from July 1976 to December 2006, moving forward in three-month steps. The final forecasts use data through September 2019 (3 months ahead) and December 2018 (12 months ahead). Crisis is defined as forecasts made using models estimated through June 2009; Post-Crisis is thereafter.

Figure 1: Conditional Distributions with Two Factors



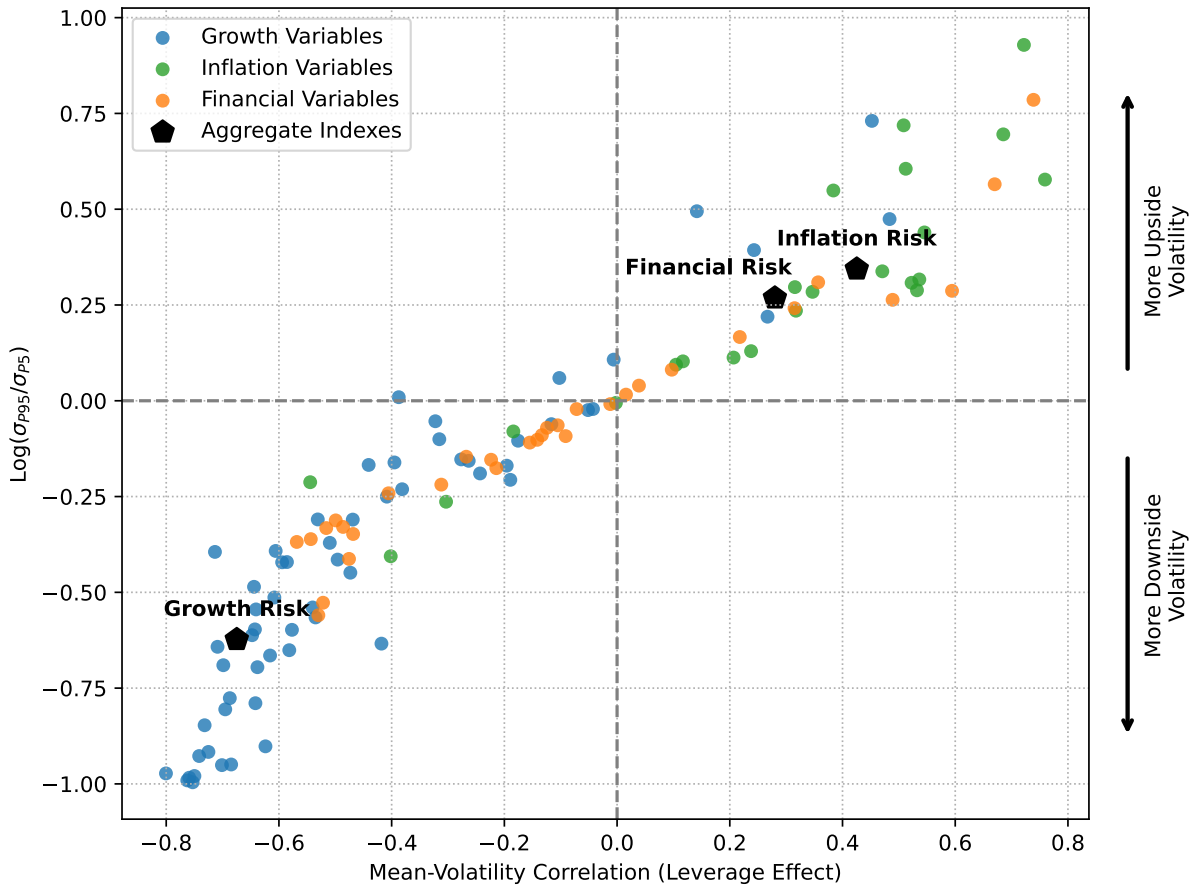
Note: This figure shows example conditional distributions for two factors (top row) and the implied distributions for three observable variables (bottom row). The conditional distributions in the bottom row are generated through linear combinations of the conditional distributions of the top row. We shut down the idiosyncratic errors when generating the conditional distributions of the observables. The vertical lines denote the 5th and 95th quantiles.

Figure 2: Aggregate Risk Indices: Growth, Inflation, and Financial Conditions



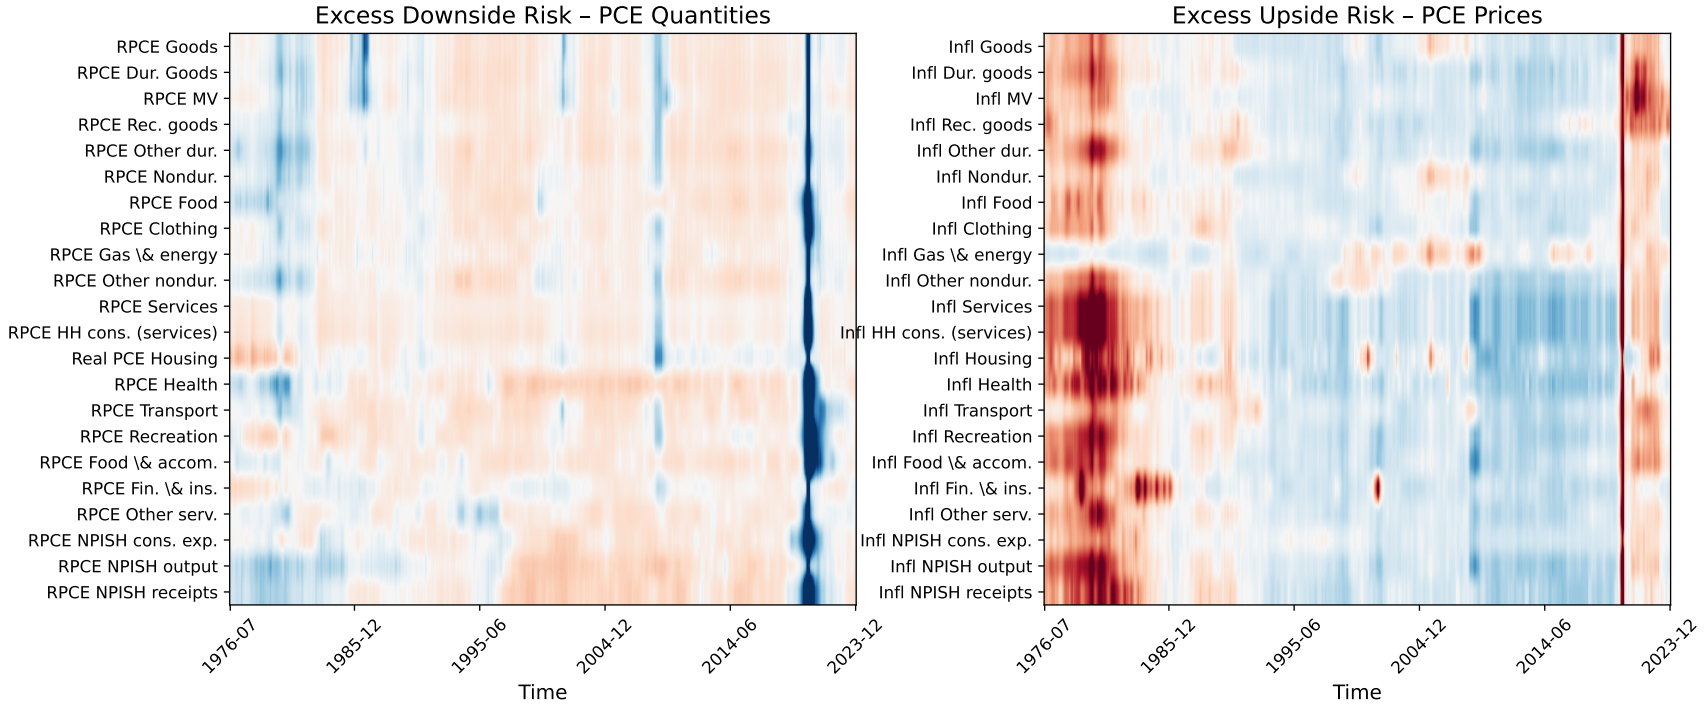
Note: This figure shows aggregate risk indices for economic growth, inflation, and financial conditions. Each series is constructed as the cross-sectional average of the $5th$ and $95th$ percentiles of the conditional distributions implied by the model for a group of related indicators. The blue solid and dashed lines report the averages of the $5th$ and $95th$ percentiles, respectively. The red solid lines report as reference the $5th$ and $95th$ percentiles of a standard normal distribution. The y-axis is truncated at ± 5.5 standard deviations to preserve readability of pre-pandemic dynamics. During the COVID-19 pandemic, peak values exceed this range: the growth risk index reaches -21.7 (5th percentile) and 16.9 (95th percentile), the financial risk index reaches -11.0 and 10.3, and the inflation risk index reaches -13.7 and 10.5. See Section 3.1 for details about the construction of the indices. Shaded areas denote NBER recessions.

Figure 3: Tail Risk Asymmetry and the Leverage Effect Across the Macroeconomy



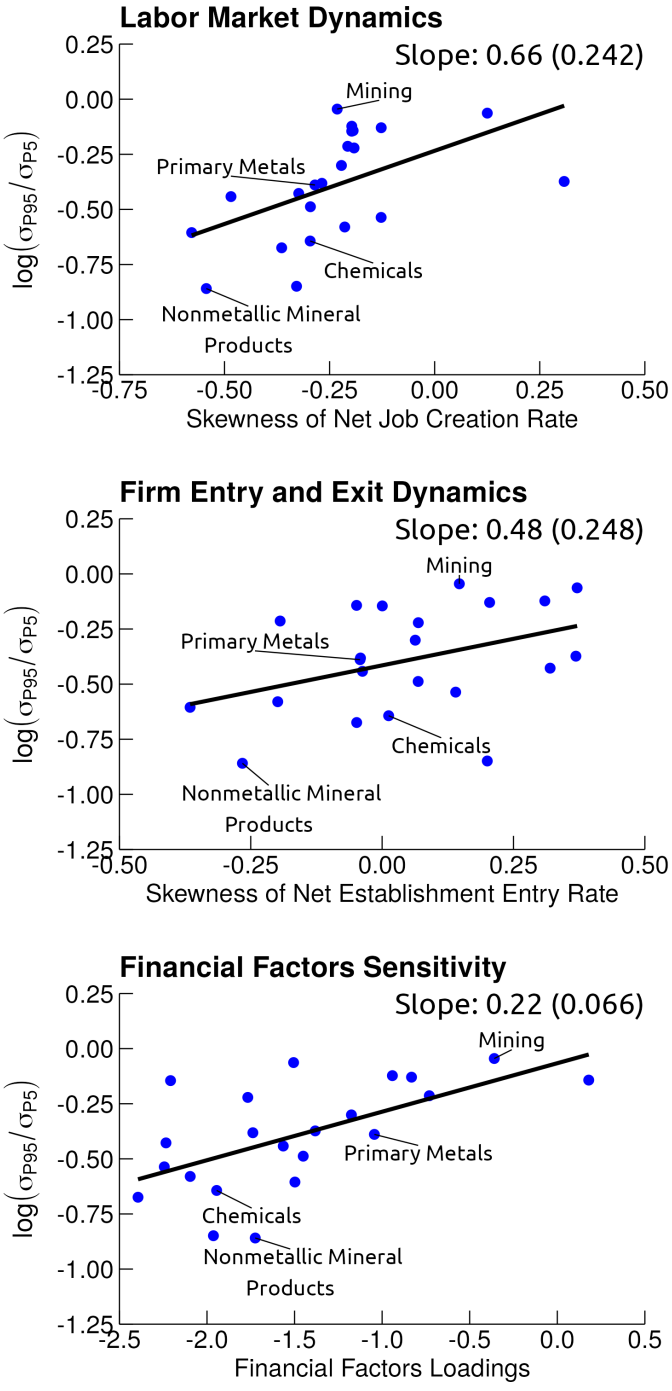
Note: For each variable, the x-axis shows the correlation between the conditional mean and volatility of its predictive distribution over time—a measure of the leverage effect—and the y-axis shows tail asymmetry, defined as the log ratio of the standard deviation of the 95th to the 5th percentile of the conditional distribution over time. Positive values indicate more volatile upper tails (upside risk); negative values indicate more volatile lower tails (downside risk). Colors denote variable categories; pentagons denote aggregate risk indices. Sample: July 1976–June 2019.

Figure 4: Heatmaps of Excessive Risk by Sector: Consumption



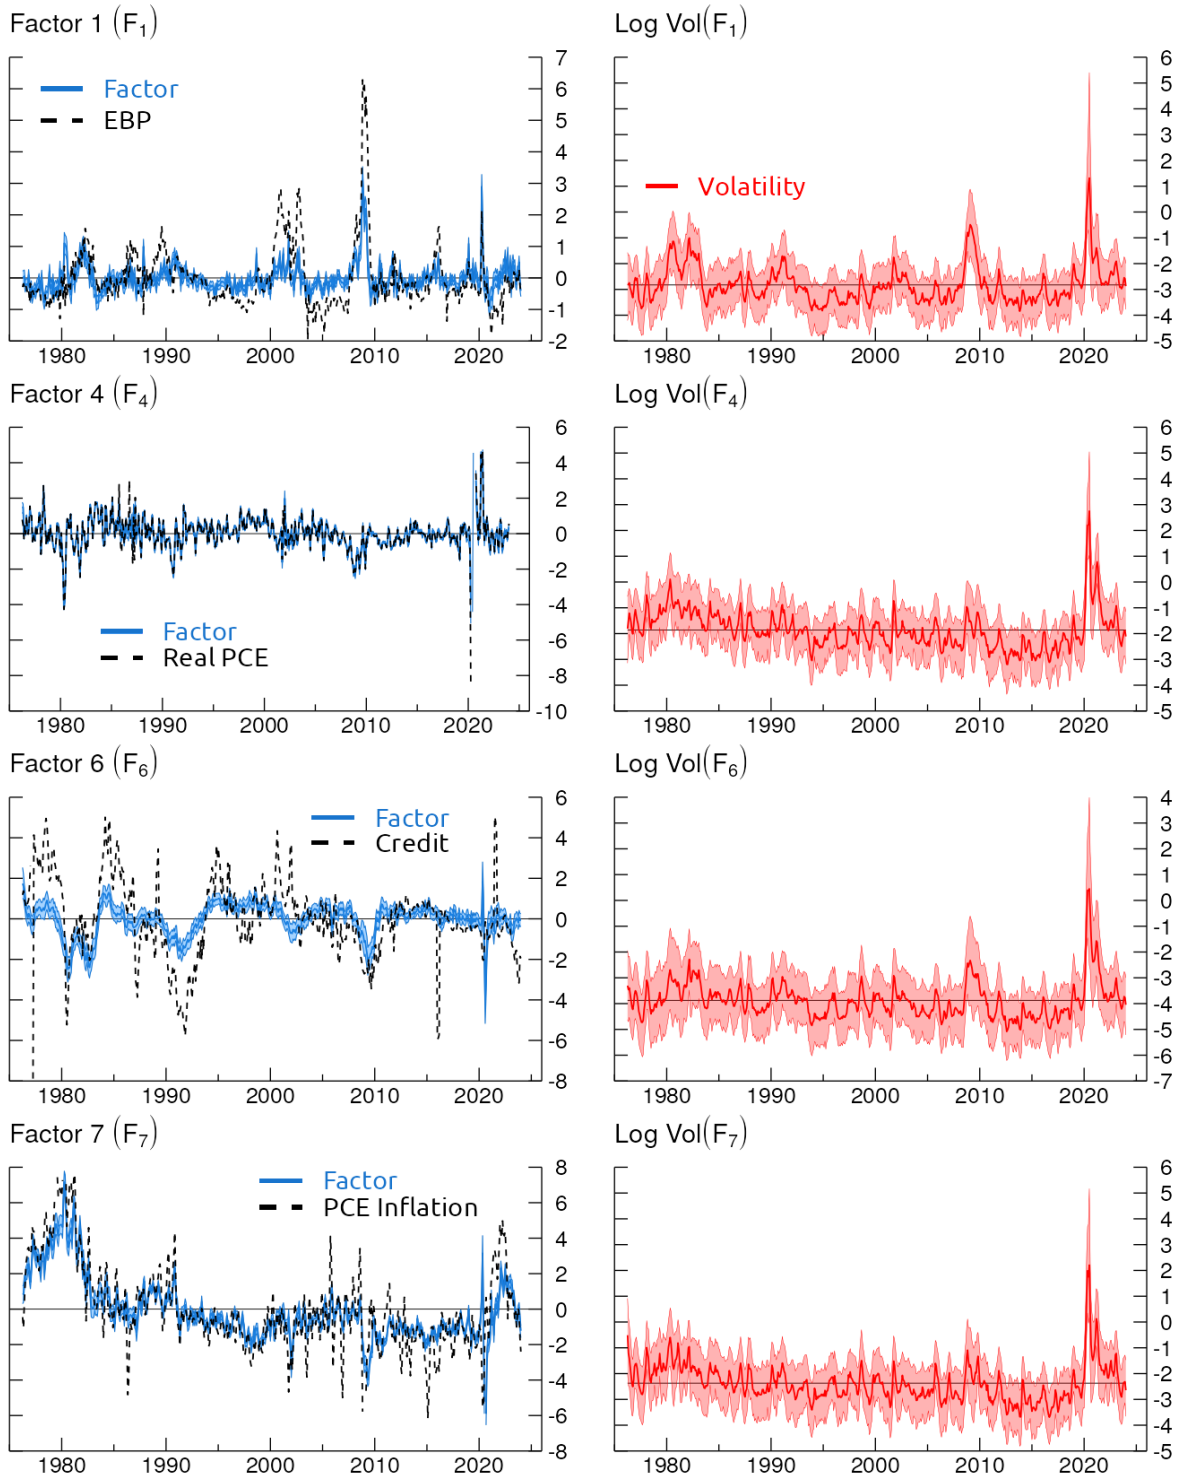
Note: The heatmaps display deviations of the conditional lower and upper quantiles of sectoral consumption expenditure from their historical averages, both quantities (left panel) and prices (right panel). The left panel shows excess downside risk, computed as the deviation of the 5th percentile from its historical mean. The right panel shows excess upside risk, based on the deviation of the 95th percentile. For quantities, blue indicates periods of elevated downside risk relative to history, while red denotes periods of subdued risk; for prices, red denotes elevated upside risk, while blue denotes periods of subdued risk.

Figure 5: Tail Risk Asymmetry and Industry Characteristics



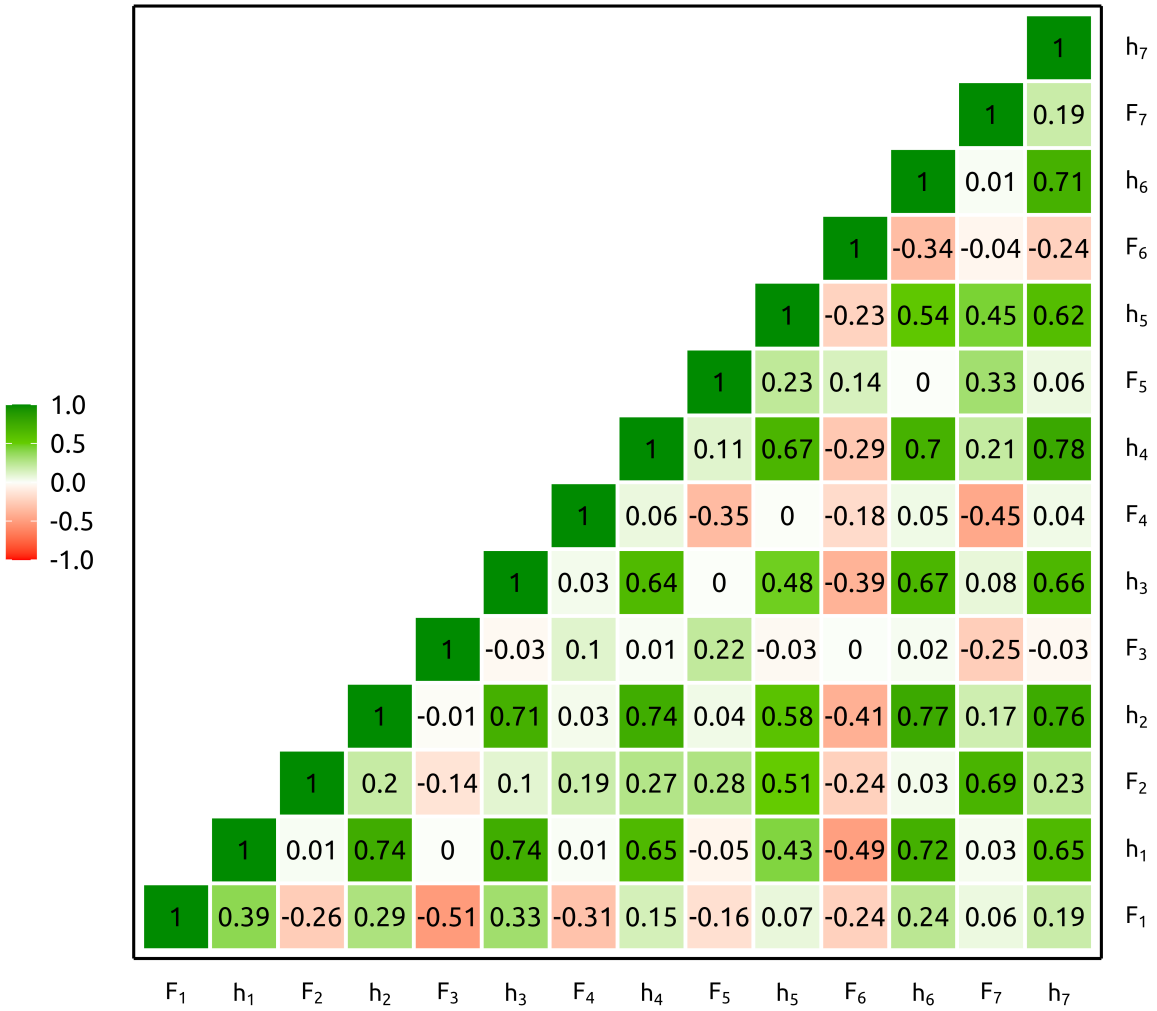
Note: This figure plots on the y-axis the log ratio of the conditional $95th$ and $5th$ percentile standard deviations for the IP sectoral variables from 1977 - 2019. On the x-axis, the panels plot the Kelley measure of skewness and the volatility for net job creation rate and net establishment entry rate by sector. The net establishment entry rate equals the establishment creation rate minus the establishment exit rate. The net job creation rate is the job creation rate minus the job destruction rate.

Figure 6: Smoothed Estimates of the Factors and Volatilities



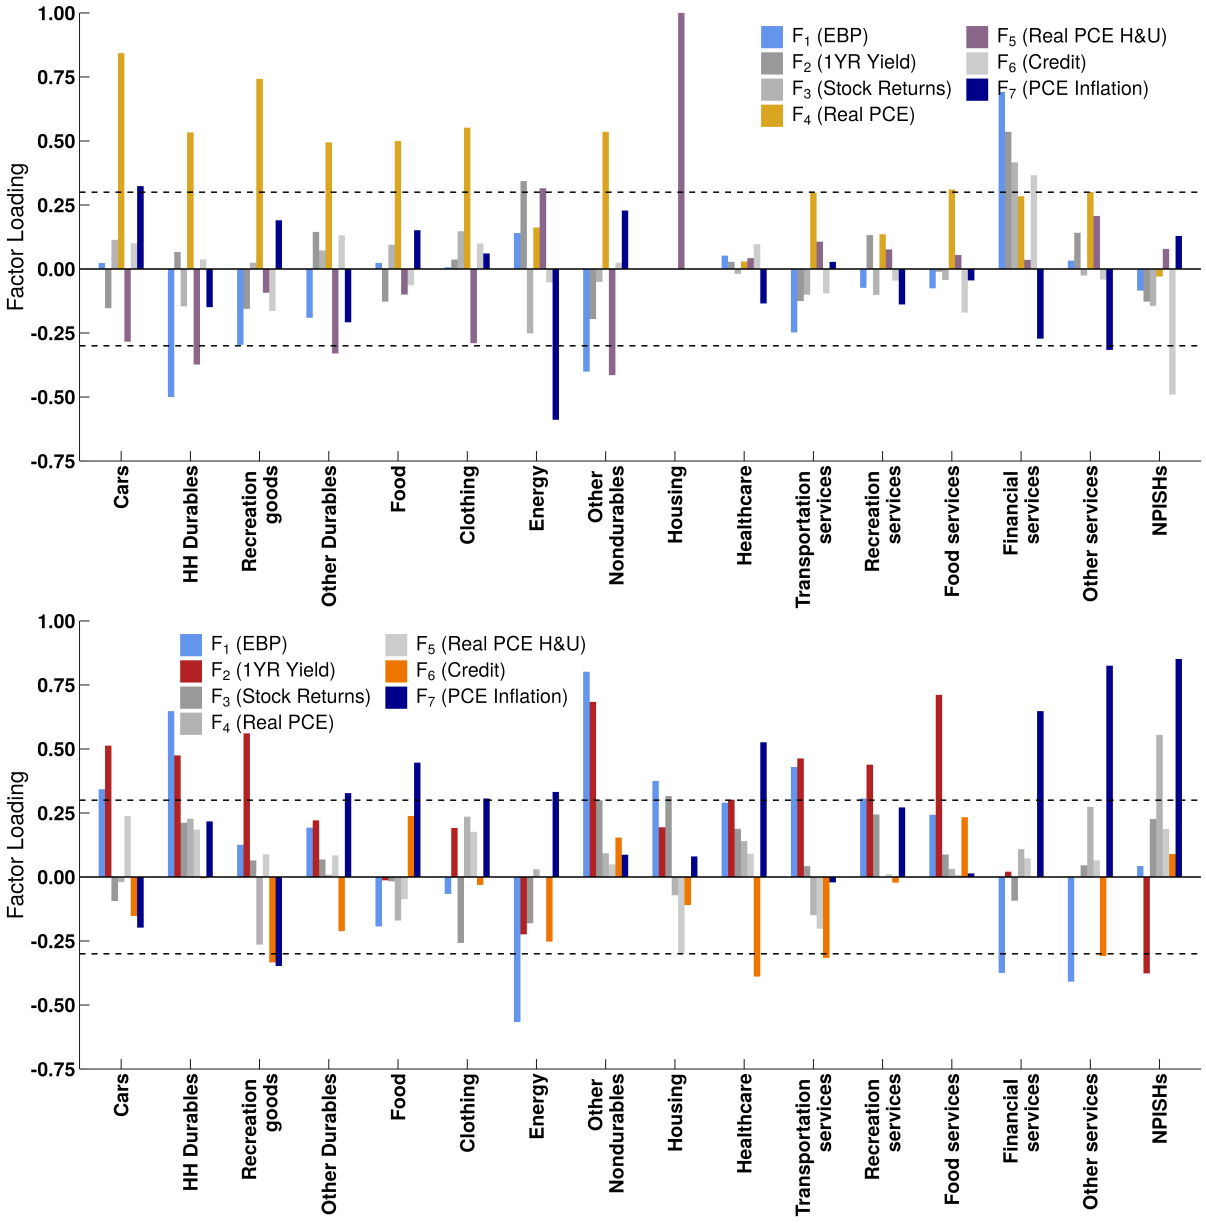
Note: The figure shows posterior median estimates of the level factors (blue lines) and of the log volatilities (red lines). Shaded areas denote the 5th and 95th of the posterior distributions. The dashed lines depict the observable variable associated to the factor for the normalization of matrix B. See Section 2 for details on the normalization.

Figure 7: Correlation across Factors



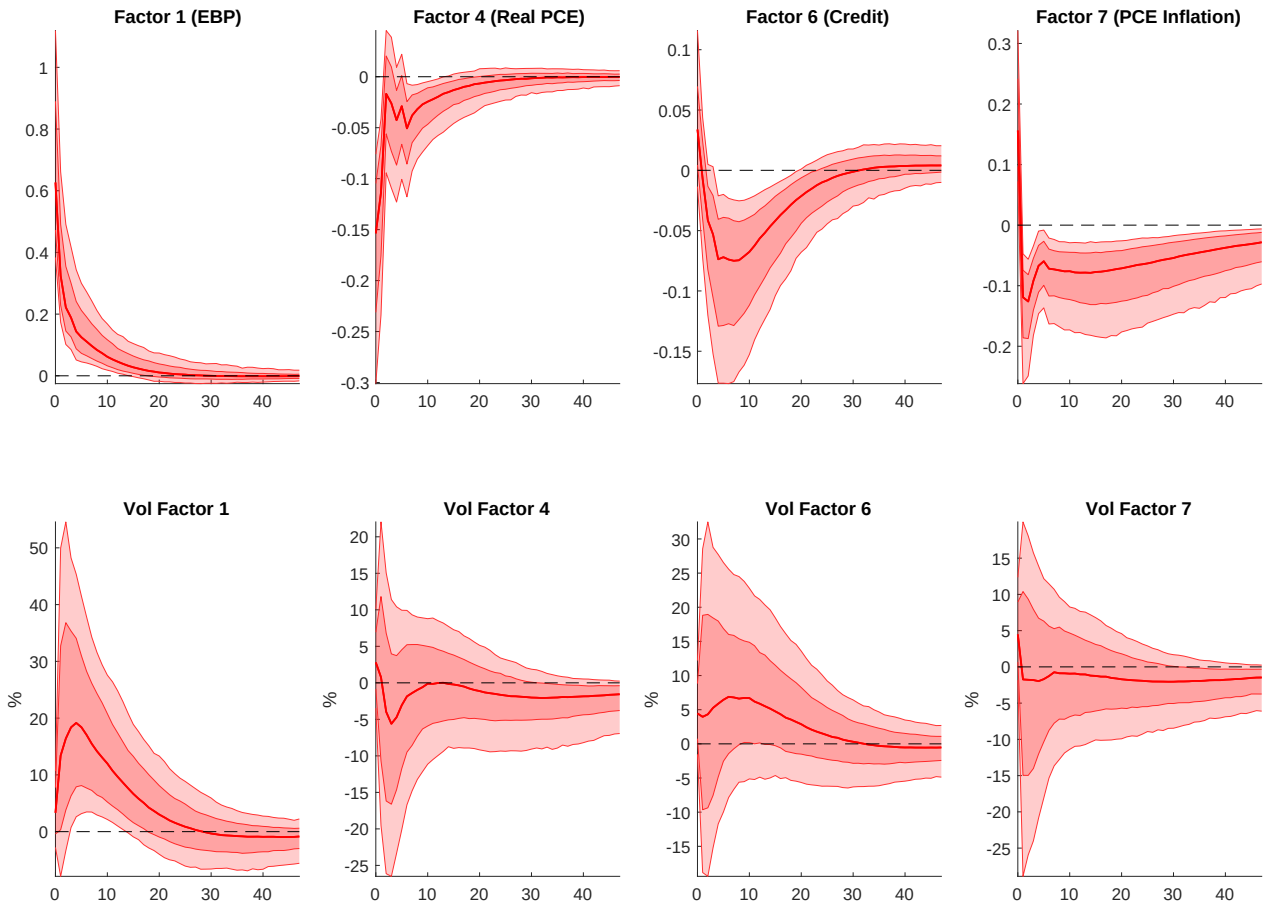
Note: This figure displays correlation coefficients among the estimated factors F_x and the log volatilities h_x.

Figure 8: Sectoral Factor Loadings for PCE Quantities (top) and Prices (bottom)



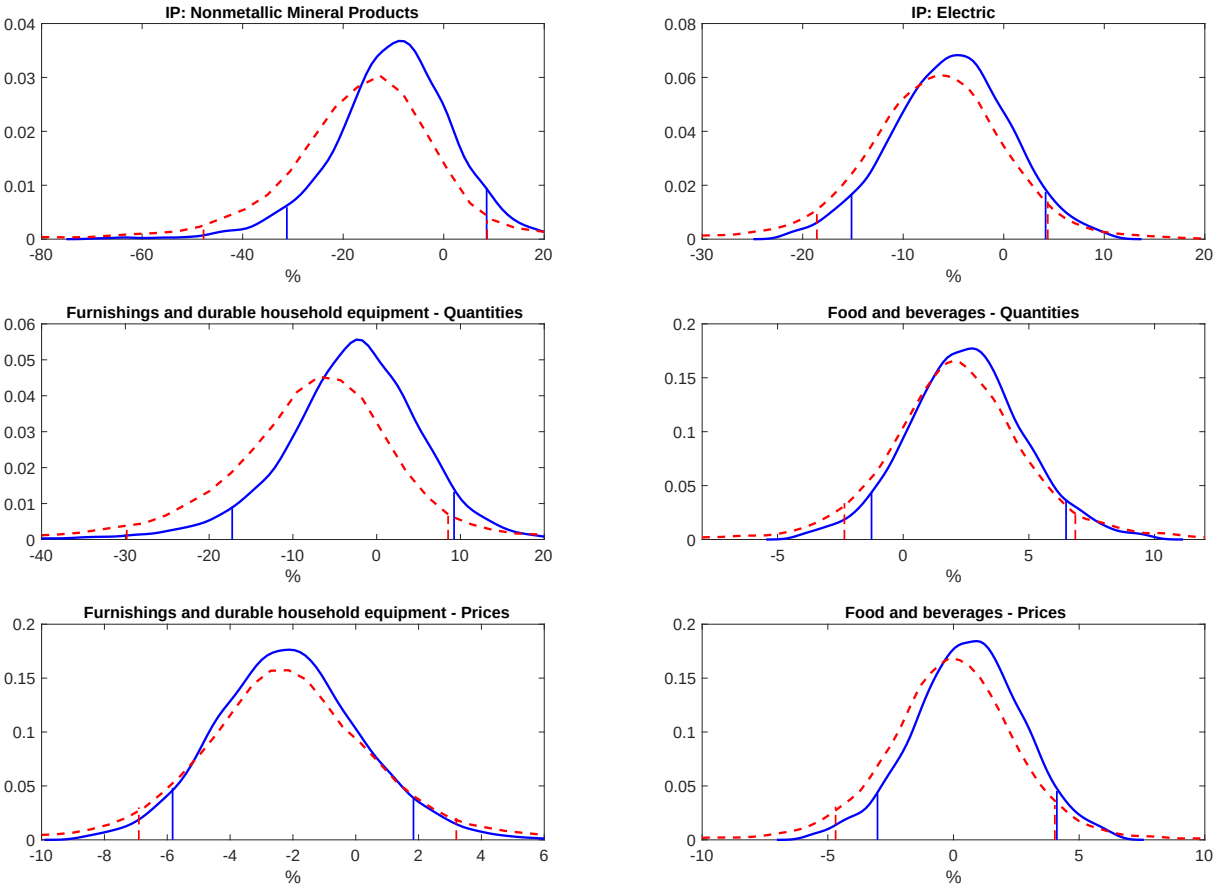
Note: This figure shows the median factor loading estimates for each PCE sector. The top panel shows the estimates for PCE Quantities and the bottom panel for PCE Prices. In the top panel, factors 1, 4, 5, and 7 are colored, representing the four most important factors for quantities. In the bottom panel, factors 1, 2, 6, and 7 are colored, representing the four most important for prices. The two dashed lines denote 0.3 and -0.3.

Figure 9: Effects of a Shock to the First Factor on Factors and Volatilities in Oct 2008



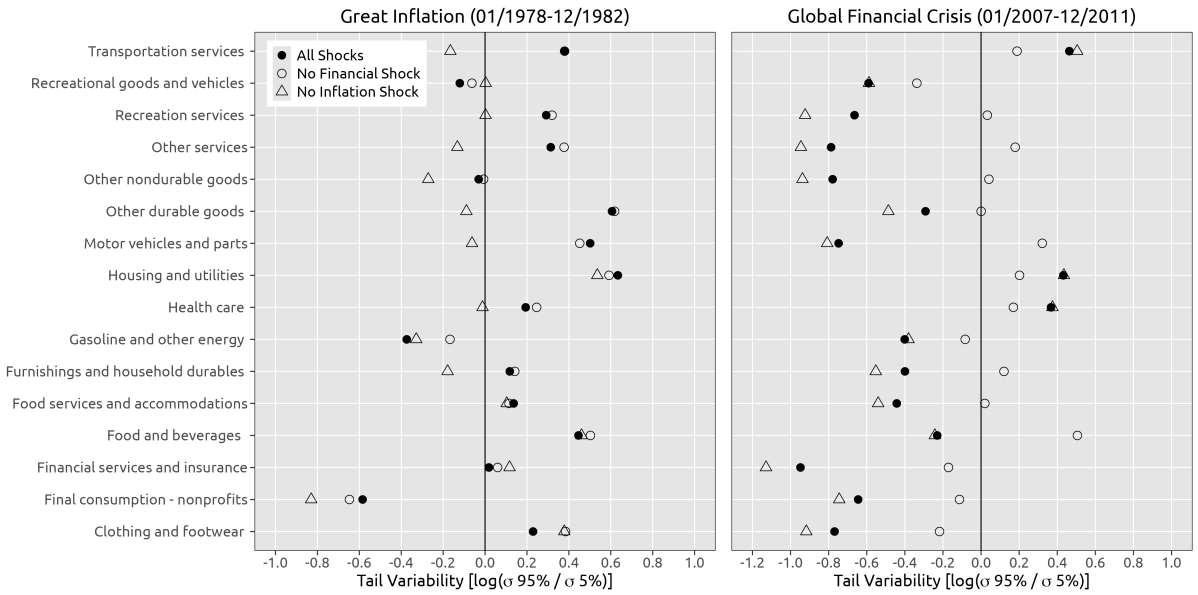
Note: This figure plots impulse response functions of a one standard deviation shock to the first factor in a Cholesky decomposition ordered first given October 2008 conditions. The first column shows the level factor responses and the second column the log volatility responses. The line is the posterior median, the dark shaded area is the 68% credible set and the light shaded area is the 90% credible set.

Figure 10: Distributional Effects of a Shock to the First Factor on IP and PCE Sectors in Oct 2008



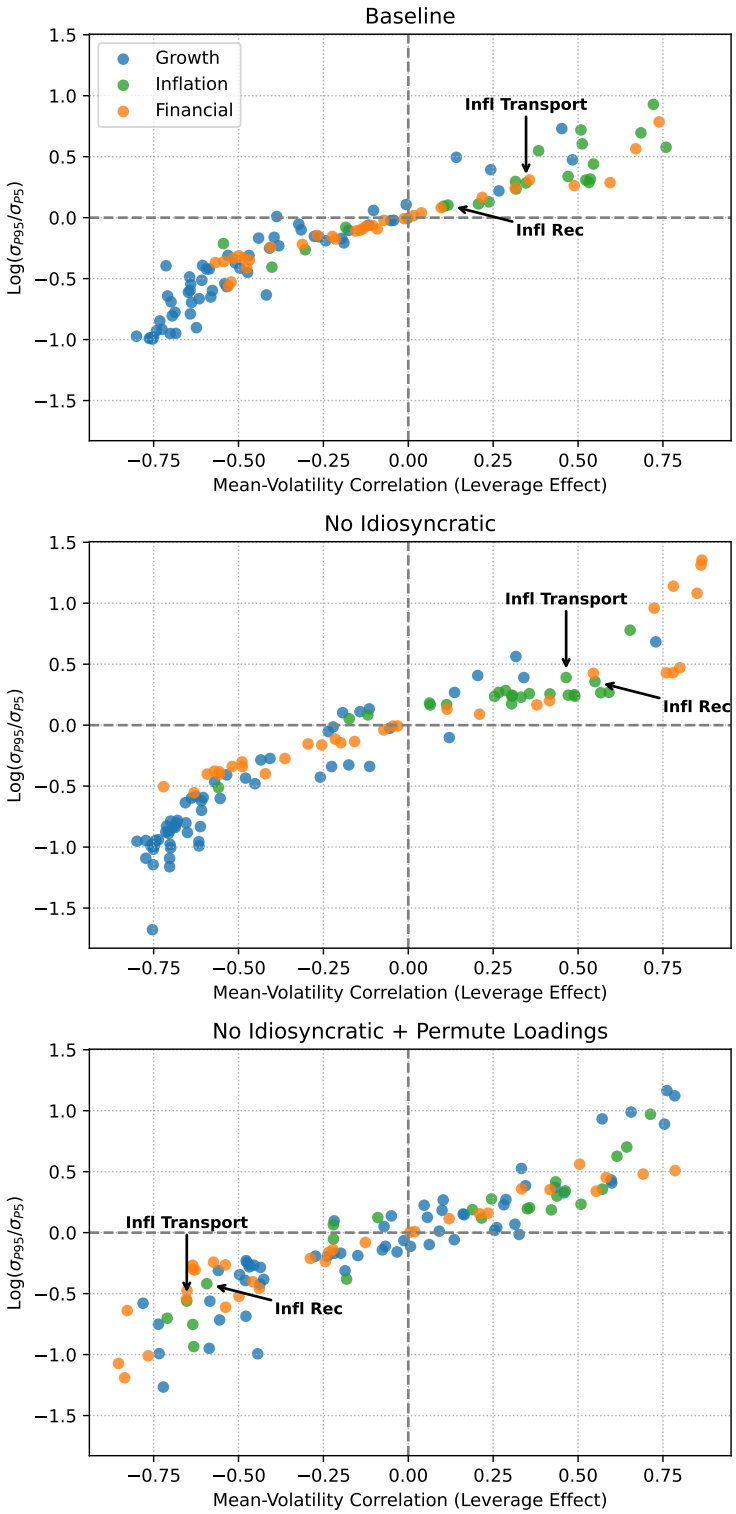
Note: This figure plots 12 month ahead distributions given October 2008 conditions. For variables in growth rates, we cumulative over the 12 months. The blue distribution is a baseline without additional shocks and the red distribution is a counterfactual in which a three standard deviation shock to the first factor in a Cholesky decomposition ordered first realizes in November 2008. The vertical lines denote the 5th and 95th quantiles.

Figure 11: Tail Variability of Price Ratios Across Expenditure Categories in the GFC (left) and Great Inflation (right)



Note: Tail variability of price ratios across different expenditure categories during two crisis periods: Great Inflation (January 1978 - December 1982) and Global Financial Crisis (January 2007 - December 2011). The x-axis shows the log ratio of the 95th to 5th percentiles. Different markers represent scenarios with all shocks (filled circles), no financial shock (open circles), and no inflation shock (triangles).

Figure 12: Mean-Volatility Correlation versus Tail Risk: Robustness



Note: The top figure reproduces the scatter in Figure 3. The second figure generates the scatter shutting off the idiosyncratic components of all variables. The third figure generates the scatter shutting off the idiosyncratic components of all variables and randomly reshuffling the factor loadings for each variable.

Online Appendix

A Data

Our dataset includes 116 monthly time series from January 1973 to December 2023, which Table A.1 summarizes. We use FRED for the aggregate real activity, financial, and inflation series, as well as the sectoral industrial production series. We use NIPA Tables 2.8.3 and 2.8.4 for the sectoral real consumption and price series. For the sectoral equity series, we use Kenneth French's database. For the excess bond premium, we use the updates provided by Favara et al. (2016) and for the corporate spread series not found in FRED, we use Global Financial Database.

We consider four modifications to the data. First, we induce stationarity to most series by either converting them to growth rates with a log difference transformation or computing a first difference. For the series that are already stationary, we either do not perform any transformation or take their logs. Second, we standardize the data using full-sample standard deviations. Third, when computing the risk indices (Figure 2) and scatters (Figures 3 and 12), we normalize each series to be in line with the discussion in Section 3.1. Finally, for robustness, we check real activity series for outliers pre-COVID and replace them with the median of the preceding five observations, following Stock and Watson (2025). Doing so does not make a difference for estimation compared to our benchmark results that do not use outlier detection.

The normalization column specifies the sign in which the series enters into the growth, financial, or inflation index. For aggregate growth variables, we classify the sign based off the sum of coefficients from a regression of the series on six leads and lags of industrial production growth, following Stock and Watson (2025). The sectoral industrial production and real PCE series enter the index with a positive sign. For the financial series, we specify that rates, spreads, and the VIX enter the index with a positive sign, while credit quantities and stock returns enter with a negative sign. Prices all enter the inflation index with a positive sign.

Table A.1: List of Data Series Used in Estimation.

Mnemonic	Full Name	Category	Transformation	Normalization
Aggregate Variables				
EBP	Excess Bond Premium	Financial	No transform	1
Y1_TSY	1-Year Treasury Rate	Financial	No transform	1
REAL_PERS_X_TR	Real personal income ex transfer receipts	Growth	Log diff	1
REAL_MANU	Real Manu. and Trade Industries Sales	Growth	Log diff	1
INIT_CLAI	Initial Claims	Growth	Log diff	-1
ALL_EMPL	All Employees: Total nonfarm	Growth	Log diff	1
HOUS_STAR	Housing Starts: Total New Privately Owned	Growth	Log	1
NEW_ORDE	New Orders for Durable Goods	Growth	Log diff	1
TOTA_BUSI	Total Business: Inventories to Sales Ratio	Growth	First diff	-1
COMM_LOANS	Commercial and Industrial Loans	Financial	Log diff	-1
REAL_ESTA	Real Estate Loans at All Commercial Banks	Financial	Log diff	-1
TOTA_NONR	Total Nonrevolving Credit	Financial	Log diff	-1
NONR_CONS	Nonrevolving consumer credit to Personal Income	Financial	First diff	-1
CONS_MOTO	Consumer Motor Vehicle Loans Outstanding	Financial	Log diff	-1
TOTA_CONS	Total Consumer Loans and Leases Outstanding	Financial	Log diff	-1
SECU_IN	Securities in Bank Credit at All Commercial Banks	Financial	Log diff	-1
EQ_SP500_COMP	S&P's Common Stock Price Index: Composite	Financial	Log diff	-1
VIX	VIX	Financial	No transform	1
MON3_TREA	3-Month Treasury Bill	Financial	No transform	1
Y10_TSY	10-Year Treasury C Minus FEDFUNDS	Financial	No transform	1
BAA_MINU	Baa - 10-Year Treasury	Financial	No transform	1
DJ_YIELD	Dow Jones Corporate Bond - 10-Year Treasury	Financial	No transform	1
HIGH10_Y	Federal Reserve 10-year High Quality Corporate Bond - 10-Year Treasury	Financial	No transform	1
BARRON_Y	Barron's Best Grade Bond Yield - 10-Year Treasury	Financial	No transform	1
IP_INDE	IP Index	Growth	Log diff	1
REAL_PERS	Real personal consumption expenditures	Growth	Log diff	1

Table A.1: Data description, continued

Mnemonic	Full Name	Category	Transformation	Normalization
PERS_CONS	PCE: Chain Index Prices	Inflation	Log diff	1
Equity				
EQ_NODUR	Equities: Consumer Nondurables	Financial	Log diff	-1
EQ_DURBL	Equities: Consumer Durables	Financial	Log diff	-1
EQ_MANUF	Equities: Manufacturing	Financial	Log diff	-1
EQ_ENRGY	Equities: Energy	Financial	Log diff	-1
EQ_HITEC	Equities: Business Equipment	Financial	Log diff	-1
EQ_TELCM	Equities: Telecommunications	Financial	Log diff	-1
EQ_SHOPS	Equities: Shops	Financial	Log diff	-1
EQ_HLTH	Equities: Healthcare	Financial	Log diff	-1
EQ_UTILS	Equities: Utilities	Financial	Log diff	-1
EQ_OTHER	Equities: Other	Financial	Log diff	-1
EQ_LO20	Equities: Size, Bottom 20%	Financial	Log diff	-1
EQ_QNT2	Equities: Size, 21%-40%	Financial	Log diff	-1
EQ_QNT3	Equities: Size, 41%-60%	Financial	Log diff	-1
EQ_QNT4	Equities: Size, 61%-80%	Financial	Log diff	-1
EQ_HI20	Equities: Size, Top 20%	Financial	Log diff	-1
Industrial Production				
IP_MANU_SIC	IP: Manufacturing (SIC)	Growth	Log diff	1
IP_MANU_NAICS	IP: Manufacturing (NAICS)	Growth	Log diff	1
IP_DURA_MAN	IP: Durable Manufacturing	Growth	Log diff	1
IP_WOOD	IP: Wood Products (NAICS = 321)	Growth	Log diff	1
IP_NONM	IP: Nonmetallic Mineral Products (NAICS = 327)	Growth	Log diff	1
IP_PRIM	IP: Primary Metals (NAICS = 331)	Growth	Log diff	1
IP_FABR	IP: Fabricated Metal Products (NAICS = 332)	Growth	Log diff	1
IP_MACH	IP: Machinery (NAICS = 333)	Growth	Log diff	1

Table A.1: Data description, continued

Mnemonic	Full Name	Category	Transformation	Normalization
IP_COMP	IP: Computer and Electronic Products (NAICS = 334)	Growth	Log diff	1
IP_ELEC_EQUIP	IP: Electrical Equip, Appliances, and Components (NAICS = 335)	Growth	Log diff	1
IP_MOTO	IP: Motor Vehicles and Parts (NAICS = 3361, 3362, 3363)	Growth	Log diff	1
IP_AERO	IP: Aerospace and Misc Transportation Equipment (NAICS = 3364-3369)	Growth	Log diff	1
IP_FURN	IP: Furniture and Related Products (NAICS = 337)	Growth	Log diff	1
IP_DURA_MISC	IP: Durable Misc (NAICS = 339)	Growth	Log diff	1
IP_NOND	IP: Nondurable Manufacturing	Growth	Log diff	1
IP_FOOD	IP: Food, Beverage, and Tobacco Products (NAICS = 311, 312)	Growth	Log diff	1
IP_TEXT	IP: Textile and Product Mills (NAICS = 313, 314)	Growth	Log diff	1
IP_APPA	IP: Apparel and Leather (NAICS = 315, 316)	Growth	Log diff	1
IP_PAPE	IP: Paper (NAICS = 322)	Growth	Log diff	1
IP_PRIN	IP: Printing and Support (NAICS = 323)	Growth	Log diff	1
IP_PETR	IP: Petroleum and Coal Products (NAICS = 324)	Growth	Log diff	1
IP_CHEM	IP: Chemicals (NAICS = 325)	Growth	Log diff	1
IP_PLAS	IP: Plastics and Rubber Products (NAICS = 326)	Growth	Log diff	1
IP_OTHE	IP: Other Manufacturing (Non-NAICS) (NAICS = 1133, 5111)	Growth	Log diff	1
IP_MINI	IP: Mining (NAICS = 21)	Growth	Log diff	1
IP_UTIL	IP: Utilities (NAICS = 2211, 2212)	Growth	Log diff	1
IP_ELEC	IP: Electric (NAICS = 2211)	Growth	Log diff	1
IP_NATU	IP: Natural Gas (NAICS = 2212)	Growth	Log diff	1
Real Personal Consumption Expenditures (Quantities)				
RPCE_GOOD_QUAN	Goods	Growth	Log diff	1
RPCE_DURA_GOOD	Durable goods	Growth	Log diff	1
RPCE_MOTO_VEHI	Motor vehicles and parts	Growth	Log diff	1
RPCE_FURN	Furnishings and durable household equipment	Growth	Log diff	1
RPCE_RECR_GOOD	Recreational goods and vehicles	Growth	Log diff	1
RPCE_OTHE_DURA	Other durable goods	Growth	Log diff	1

Table A.1: Data description, continued

Mnemonic	Full Name	Category	Transformation	Normalization
RPCE_NOND_GOOD	Nondurable goods	Growth	Log diff	1
RPCE_FOOD	Food and beverages purchased for off-premises consumption	Growth	Log diff	1
RPCE_CLOT	Clothing and footwear	Growth	Log diff	1
RPCE_GASO	Gasoline and other energy goods	Growth	Log diff	1
RPCE_OTHE_NOND	Other nondurable goods	Growth	Log diff	1
RPCE_SERV_QUAN	Services	Growth	Log diff	1
RPCE_HOUS_CONS	Household consumption expenditures (for services)	Growth	Log diff	1
RPCE_HOUS	Housing and utilities	Growth	Log diff	1
RPCE_HEAL_CARE	Health care	Growth	Log diff	1
RPCE_TRAN_SERV	Transportation services	Growth	Log diff	1
RPCE_RECR_SERV	Recreation services	Growth	Log diff	1
RPCE_FOOD_SERV	Food services and accommodations	Growth	Log diff	1
RPCE_FINA_SERV	Financial services and insurance	Growth	Log diff	1
RPCE_OTHE_SERV	Other services	Growth	Log diff	1
RPCE_FINA_CONS	Final consumption expenditures of NPISHs	Growth	Log diff	1
RPCE_GROS_OUTP	Gross output of nonprofit institutions	Growth	Log diff	1
RPCE_LESS_RECE	Less: Receipts from sales of goods and services by nonprofit inst	Growth	Log diff	1
Personal Consumption Expenditure Price Indices				
P_GOOD_PRIC	Goods	Inflation	Log diff	1
P_DURA_GOOD	Durable goods	Inflation	Log diff	1
P_MOTO_VEHI	Motor vehicles and parts	Inflation	Log diff	1
P_FURN_	Furnishings and durable household equipment	Inflation	Log diff	1
P_RECR_GOOD	Recreational goods and vehicles	Inflation	Log diff	1
P_OTHE_DURA	Other durable goods	Inflation	Log diff	1
P_NOND_GOOD	Nondurable goods	Inflation	Log diff	1
P_FOOD	Food and beverages purchased for off-premises consumption	Inflation	Log diff	1
P_CLOT	Clothing and footwear	Inflation	Log diff	1

Table A.1: Data description, continued

Mnemonic	Full Name	Category	Transformation	Normalization
P_GASO	Gasoline and other energy goods	Inflation	Log diff	1
P_OTHE_NOND	Other nondurable goods	Inflation	Log diff	1
P_SERV_PRIC	Services	Inflation	Log diff	1
P_HOUS_CONS	Household consumption expenditures (for services)	Inflation	Log diff	1
P_HOUS	Housing and utilities	Inflation	Log diff	1
P_HEAL_CARE	Health care	Inflation	Log diff	1
P_TRAN_SERV	Transportation services	Inflation	Log diff	1
P_RECR_SERV	Recreation services	Inflation	Log diff	1
P_FOOD_SERV	Food services and accommodations	Inflation	Log diff	1
P_FINA_SERV	Financial services and insurance	Inflation	Log diff	1
P_OTHE_SERV	Other services	Inflation	Log diff	1
P_FINA_CONS	Final consumption expenditures of NPISHs	Inflation	Log diff	1
P_GROS_OUTP	Gross output of nonprofit institutions	Inflation	Log diff	1
P_LESS_RECE	Less: Receipts from sales of goods and services by nonprofit inst	Inflation	Log diff	1

Note: Log diff (diff) refers to the first difference of the log (first difference). Normalization indicates the sign of the loading on the aggregate risk indices.

B Model Estimation

This appendix provides technical details on the estimation of the dynamic factor model with endogenous stochastic volatility presented in Section 2. The model is specified by equations 1–4 in the main text. We describe the prior distributions, the Gibbs sampling algorithm, and provide Monte Carlo evidence that the algorithm accurately recovers model parameters.

B.1 Prior Distributions

We employ standard conjugate priors where possible and non-informative priors otherwise. Table A.2 summarizes the prior specifications. Prior hyperparameters are calibrated using the first 30 observations (training sample) when needed.

Table A.2: Prior Specifications

Parameter	Prior Distribution	Hyperparameters
Observation Equation		
Factor loadings B_i (row i)	$N(B_0, V_{B0})$	B_0: OLS, $V_{B0} = 10 \times$ OLS var
AR coefficients $\tilde{\rho}_i$	$N(0, 0.1 \times I_L)$	L: number of lags
Idiosyncratic vol. variance g_i	$IG(0.001, 1)$	Scale $= 0.001$, df $= 1$
Transition Equation		
VAR coefficients Γ, $\tilde{\Gamma}$	$N(\Gamma_0, P_0)$	Via dummy observations (see below)
Volatility shock variance S	$IG(0.001, 1)$	Scale $= 0.001$, df $= 1$
Correlation matrix Σ	Via Huang-Wand (2013)	$a_{kj} \sim N(0, 1)$
Initial Conditions		
Factors F_0	$N(\hat{F}_1, I_N)$	$\hat{F}_1$: first principal component
Log volatilities h_{i0}	$N(\mu_{0,i}, 1)$	$\mu_{0,i}$: from training sample VAR

VAR Coefficient Priors via Dummy Observations

The priors for VAR coefficients $\Gamma = \text{vec}([c; \beta_j; b_k])$ in equation 3 and $\tilde{\Gamma} = \text{vec}([\alpha; \theta; d_j])$ in equation 4 are implemented through dummy observations, following Banbura et al. (2010). This approach allows us to specify Minnesota-style priors in a straightforward manner.

The dummy observations are:

$$y_D = \begin{bmatrix} \frac{\text{diag}(\gamma_1 s_1, \ldots, \gamma_N s_N)}{\tau} \\ 0_{N(P-1)\times N} \\ 0_{EX\times N} \end{bmatrix}, \quad x_D = \begin{bmatrix} \frac{J_P \otimes \text{diag}(s_1, \ldots, s_N)}{\tau} & 0_{NP\times EX} \\ 0_{N\times NP} & 0_{N\times EX} \\ 0_{EX\times NP} & I_{EX}/c \end{bmatrix} \tag{1}$$

where γ_i denotes the prior mean for the first own-lag coefficient (obtained from univariate AR(1) regressions), s_i is the residual standard deviation from the AR(1) regression, τ controls overall tightness, c controls tightness on exogenous regressors, EX is the number of exogenous variables, and $J_P = \text{diag}(1, 2, \ldots, P)$ implements lag decay.

We set $\tau = 0.1$ (tight prior on VAR dynamics). For coefficients on lagged volatilities, we use $c = 0.1$; for intercepts, we use $c = 1000$ (flat prior). The prior moments are then:

$$\Gamma_0 = (x_D' x_D)^{-1}(x_D' y_D), \tag{2}$$

$$P_0 = S \otimes (x_D' x_D)^{-1}, \tag{3}$$

where S is a diagonal matrix with training sample estimates of factor variances on the diagonal.

Initial Conditions

We initialize the factors at $F_0 \sim N(\hat{F}_1, I_N)$ where $\hat{F}_1$ denotes the first observation of the principal component estimate of the factors. For the log volatilities, we estimate a VAR on the training sample, compute the Cholesky decomposition of the residual covariance matrix, and set $h_{i0} \sim N(\mu_{0,i}, 1)$ where $\mu_{0,i}$ are the log of the Cholesky diagonal elements.

B.2 Gibbs Sampling Algorithm

We draw from the posterior distribution using a Gibbs sampler that cycles through seven blocks. Let Θ denote all model parameters. In each step below, we draw from the conditional posterior given the most recent draws of all other parameters.

Step 1: Latent Factors, $G(F_t|\Theta)$

Conditional on the stochastic volatilities $\tilde{h}_t$, r_t and other parameters, the model has a linear Gaussian state-space representation. We stack the system as:

$$Z_t = \Lambda + \Phi Z_{t-1} + E_t, \quad \text{var}(E_t) = Q_{F,t}, \tag{4}$$

$$\tilde{X}_t = \tilde{B} Z_t + V_t, \quad \text{var}(V_t) = R_{F,t}, \tag{5}$$

where the state vector is $Z_t = [\tilde{h}'_{t+1}, F'_t, \tilde{h}'_t, F'_{t-1}, \ldots]'$. The matrices Φ and Λ capture the VAR dynamics with appropriate lags, $\tilde{X}_t$ contains the AR-filtered observables, and $\tilde{B}$ adjusts for the AR structure in idiosyncratic shocks. The covariance matrix $Q_{F,t}$ contains Ω_t (equation 6) in the upper-left block with zeros elsewhere, and $R_{F,t} = \text{diag}(0_N, R_t)$ reflects time-varying idiosyncratic volatility.

We apply the Carter and Kohn (1994) algorithm to draw the state vector from its conditional posterior distribution, which is Gaussian given the linear structure and known volatilities.

Step 2: Factor Loadings B **and AR Coefficients** ρ_i

Given factors F_t and volatilities R_t, the observation equation becomes n independent regressions with known heteroscedasticity. For row i of B:

$$X^*_{it} = B_i F^*_t + \tilde{v}_{it}, \quad \tilde{v}_{it} \sim N(0,1), \tag{6}$$

where $X^*_{it} = (X_{it} - \sum_{l=1}^{L} \rho_{i,l} X_{it-l})/\sqrt{\exp(r_{it})}$ and F^*_t is similarly transformed. The conditional posterior is $N(m_i, v_i)$ with:

$$v_i = (V_{B0}^{-1} + F_t^{*'} F_t^*)^{-1}, \quad m_i = v_i (V_{B0}^{-1} B_0 + F_t^{*'} X^*_{it}). \tag{7}$$

The AR coefficients $\tilde{\rho}_i = [\rho_{i,1}, \ldots, \rho_{i,L}]$ are drawn similarly after computing residuals $v_{it} = X_{it} - B_i F_t$ and GLS-transforming by $\sqrt{\exp(r_{it})}$.

Step 3: Idiosyncratic Volatilities, $G(r_{it}|\Theta)$

Conditional on factors and factor loadings, we obtain n independent univariate stochastic volatility models:

$$u_{it} = \sqrt{\exp(r_{it})}\bar{u}_{it}, \quad \bar{u}_{it} \sim N(0,1), \tag{8}$$

$$r_{it} = r_{it-1} + g_i^{1/2}\epsilon_{it}, \quad \epsilon_{it} \sim N(0,1). \tag{9}$$

We draw r_{it} using the algorithm of Jacquier et al. (2002). Given r_{it}, the variance parameter g_i has an inverse Gamma posterior: $IG((r_{it} - r_{it-1})'(r_{it} - r_{it-1}) + 0.001, T + 1)$.

Step 4: VAR Coefficients, $G(\Gamma, \tilde{\Gamma}|\Theta)$

Conditional on F_t, S, H_t, and Σ, equations 3 and 4 form a SUR system with time-varying heteroscedasticity:

$$Y_t = X_t\Pi_t + E_t, \quad \text{var}(E_t) = G_t\Sigma G_t', \tag{10}$$

where $G_t = \text{diag}([\tilde{s}^{1/2}, \exp(\tilde{h}_t)^{1/2}])$ and $\Pi_t = [\Gamma', \tilde{\Gamma}']'$ stacks the coefficients. Given the time-varying covariance structure, we use the Kalman filter to compute the posterior distribution $N(\Pi_{T|T}, P_{T|T})$ with transition equation $\Pi_t = \Pi_{t-1}$. The Kalman recursions are:

$$\Pi_{t|t-1} = \Pi_{t-1|t-1}, \quad P_{t|t-1} = P_{t-1|t-1}, \tag{11}$$

$$\eta_{t|t-1} = Y_t - X_t\Pi_{t|t-1}, \quad f_{t|t-1} = X_tP_{t|t-1}X_t' + G_t\Sigma G_t', \tag{12}$$

$$K_t = P_{t|t-1}X_t'f_{t|t-1}^{-1}, \tag{13}$$

$$\Pi_{t|t} = \Pi_{t|t-1} + K_t\eta_{t|t-1}, \quad P_{t|t} = P_{t|t-1} - K_tX_tP_{t|t-1}. \tag{14}$$

The filter is initialized at Π_0 and P_0 from the prior.

Step 5: Variance of Volatility Shocks, $G(S|\Theta)$

Given the residuals η_t from equation 4 and Σ, the transition equation can be written as:

$$\tilde{h}_t - S^{1/2}\mu_{\eta_t|e_t} = \alpha + \theta\tilde{h}_{t-1} + \sum_{j=1}^{Q} d_jF_{t-j} + \eta_t^*, \quad \text{var}(\eta_t^*) = S^{1/2}\Sigma_{\eta|e}S^{1/2'}, \tag{15}$$

where $\mu_{\eta_t|e_t} = e_t\Sigma_e^{-1}\Sigma_{\eta e}$ and $\Sigma_{\eta|e} = \Sigma_\eta - \Sigma_{\eta e}'\Sigma_e^{-1}\Sigma_{\eta e}$.

The correlation between volatility shocks makes the conditional posterior non-standard,

requiring a Metropolis step. We use a mixture proposal:

$$q(S_j) = \varkappa \cdot IG(v_1, T_1) + (1 - \varkappa) \cdot IG(v(S_{j-1}), T(\bar{V})), \tag{16}$$

where $v_1 = \tilde{\eta}_{it}'\tilde{\eta}_{it} + 0.001$ and $T_1 = T + 1$ (based on residuals $\tilde{\eta}_{it}$ from the ith equation). The second component is centered at the previous draw with standard deviation $\bar{V}$:

$$v(S_{j-1}) = 2S_{j-1}\left(1 + \frac{S_{j-1}^2}{\bar{V}^2}\right), \quad T(\bar{V}) = 2\left(2 + \frac{S_{j-1}^2}{\bar{V}^2}\right). \tag{17}$$

We set $\varkappa = 0.5$ and tune $\bar{V}$ for an acceptance rate near 25-40%. The draw is accepted with probability:

$$\alpha = \min\left\{1, \frac{g(E_t|S_j)q(S_{j-1})}{g(E_t|S_{j-1})q(S_j)}\right\}, \tag{18}$$

where $g(E_t|S_j)$ is the likelihood of the transition equation residuals given S_j.

Step 6: Covariance Matrix, $G(\Sigma|\Theta)$

We draw the restricted covariance matrix Σ (with unit diagonal) using the independence Metropolis algorithm of Chan and Jeliazkov (2009). They decompose $\Sigma = L^{-1}DL^{-1'}$ where L is lower triangular with ones on the diagonal and D is diagonal. The unit diagonal restriction on Σ implies:

$$\lambda_1 = 1, \tag{19}$$

$$\lambda_k = 1 - \sum_{j=1}^{k-1}(a_{kj})^2\lambda_j, \quad k = 2, \ldots, N + n, \tag{20}$$

where a_{kj} are the lower-triangular elements of L^{-1} and λ_k are the diagonal elements of D.

The proposal density for the vector $\mathbf{a} = [a_{21}, a_{31}, a_{32}, \ldots]'$ is:

$$f(\mathbf{a}|\varepsilon_t) = N(\mu, \tau V), \tag{21}$$

where $V = (A_0^{-1} + \sum_{t=1}^{T} U_t\hat{D}^{-1}U_t)^{-1}$, $\mu = V(A_0^{-1}\mathbf{a}_0 + \sum_{t=1}^{T} U_t\hat{D}^{-1}\varepsilon_t)$, and U_t is a matrix constructed from the residuals ε_t as described in Chan and Jeliazkov (2009). The diagonal matrix $\hat{D}$ is obtained by iterating between the formula for μ and the restriction equations for λ_k. The draw is accepted with probability:

$$\alpha = \min\left\{1, \frac{g(\varepsilon_t|\Sigma_{\text{new}})f(\mathbf{a}_{\text{old}}|\varepsilon_t)}{g(\varepsilon_t|\Sigma_{\text{old}})f(\mathbf{a}_{\text{new}}|\varepsilon_t)}\right\}, \tag{22}$$

with the restriction that $\lambda_k > 0$ for all k to ensure positive definiteness.

Step 7: Stochastic Volatilities, $G(\tilde{h}_t|\Theta)$

This is the most computationally intensive step. Conditional on factors and other parameters, we have a multivariate non-linear state-space system. We rewrite it as:

$$\boldsymbol{F}_t = C + \Psi \boldsymbol{F}_{t-1} + N_t, \tag{23}$$

$$F_t - H_t^{1/2}\mu_{e_t|\eta_t} = c + \sum_{j=1}^{P} \beta_j F_{t-j} + \sum_{k=1}^{K} b_k \tilde{h}_{t-k} + \tilde{e}_t, \tag{24}$$

where $\boldsymbol{F}_t = [\eta'_{t+1}, \eta'_t, \tilde{h}'_t, \tilde{h}'_{t-1}, \ldots]'$ includes the volatility shocks as state variables, $\mu_{e_t|\eta_t} = \eta_t \Sigma_\eta^{-1} \Sigma'_{\eta e}$ is the conditional mean, $\Sigma_{e_t|\eta_t} = \Sigma_e - \Sigma_{\eta e}\Sigma_\eta^{-1}\Sigma'_{\eta e}$ is the conditional variance, and $\text{var}(\tilde{e}_t) = H_t^{1/2}\Sigma_{e|\eta}H_t^{1/2'}$. Importantly, $\tilde{e}_t$ is uncorrelated with N_t.

We use particle Gibbs with ancestor sampling (Andrieu et al., 2010; Lindsten et al., 2014) to sample from the conditional posterior. Let $\boldsymbol{F}_t^{(i-1)}$ denote the trajectory from the previous Gibbs iteration. The algorithm with $\tilde{M}$ particles proceeds as:

For $t = 1$**:**

1. Draw $\boldsymbol{F}_1^{(j)}$ for $j = 1, \ldots, \tilde{M} - 1$. Fix $\boldsymbol{F}_1^{(\tilde{M})} = \boldsymbol{F}_1^{(i-1)}$.
2. Compute normalized weights $p_1^{(j)} = w_1^{(j)} / \sum_{j=1}^{\tilde{M}} w_1^{(j)}$ where $w_1^{(j)}$ is the likelihood:

$$w_1^{(j)} = |\Omega_1^{(j)}|^{-1/2} \exp\left(-\frac{1}{2}\tilde{e}'_1(\Omega_1^{(j)})^{-1}\tilde{e}_1\right). \tag{25}$$

For $t = 2$ **to** T**:**

1. Resample indices $a_t^{(j)}$ for $j = 1, \ldots, \tilde{M} - 1$ with $P(a_t^{(j)} = k) \propto p_{t-1}^{(k)}$.
2. Draw $\boldsymbol{F}_t^{(j)}|\boldsymbol{F}_{t-1}^{(a_t^{(j)})}$ from the transition equation for $j = 1, \ldots, \tilde{M} - 1$.
3. Fix $\boldsymbol{F}_t^{(\tilde{M})} = \boldsymbol{F}_t^{(i-1)}$.
4. **Ancestor sampling:** Draw $a_t^{(\tilde{M})}$ with probability proportional to:

$$w_{t-1}^{(j)} \prod_{s=t}^{t-1+\mathcal{L}} g(F_s|\boldsymbol{F}_{1:t-1}^{(j)}, \boldsymbol{F}_{t:s}^{(i-1)}) f(\boldsymbol{F}_s^{(i-1)}|\boldsymbol{F}_{1:t-1}^{(j)}, \boldsymbol{F}_{t:s-1}^{(i-1)}), \tag{26}$$

where $\mathcal{L}$ is a look-ahead parameter (we set $\mathcal{L} = 5$). This step breaks the reference trajectory into pieces, preventing particle degeneracy.

5. Update weights as in step (2) above.

Final step: Draw $F_t^{(i)}$ by sampling index j with probability $p_T^{(j)}$.

We use $\tilde{M} = 20$ particles. The ancestor sampling step is critical for good mixing; without it, the particle system collapses to the reference trajectory.

B.3 Monte Carlo Validation

To verify the algorithm's performance, we conduct a Monte Carlo experiment with a simplified two-factor version of the model. We generate 600 observations, discarding the first 100 as burn-in and retaining 500 for analysis, from the following data generating process:

$$\begin{pmatrix} \ln h_{1,t+1} \\ \ln h_{2,t+1} \end{pmatrix} = \begin{pmatrix} 0.97 & -0.01 \\ 0.01 & 0.97 \end{pmatrix} \begin{pmatrix} \ln h_{1t} \\ \ln h_{2t} \end{pmatrix} + \begin{pmatrix} -0.05 & 0.01 \\ -0.05 & 0.01 \end{pmatrix} \begin{pmatrix} F_{1,t-1} \\ F_{2,t-1} \end{pmatrix} + \begin{pmatrix} 0.1^{1/2} e_{1t} \\ 0.1^{1/2} e_{2t} \end{pmatrix}, \tag{27}$$

$$\begin{pmatrix} F_{1t} \\ F_{2t} \end{pmatrix} = \begin{pmatrix} 0.3 \\ -0.3 \end{pmatrix} + \begin{pmatrix} 0.75 & -0.1 \\ 0.1 & 0.75 \end{pmatrix} \begin{pmatrix} F_{1,t-1} \\ F_{2,t-1} \end{pmatrix} + \begin{pmatrix} -0.05 & 0.01 \\ -0.05 & 0.01 \end{pmatrix} \begin{pmatrix} \ln h_{1,t-1} \\ \ln h_{2,t-1} \end{pmatrix} + \begin{pmatrix} h_{1t}^{1/2} e_{3t} \\ h_{2t}^{1/2} e_{4t} \end{pmatrix}, \tag{28}$$

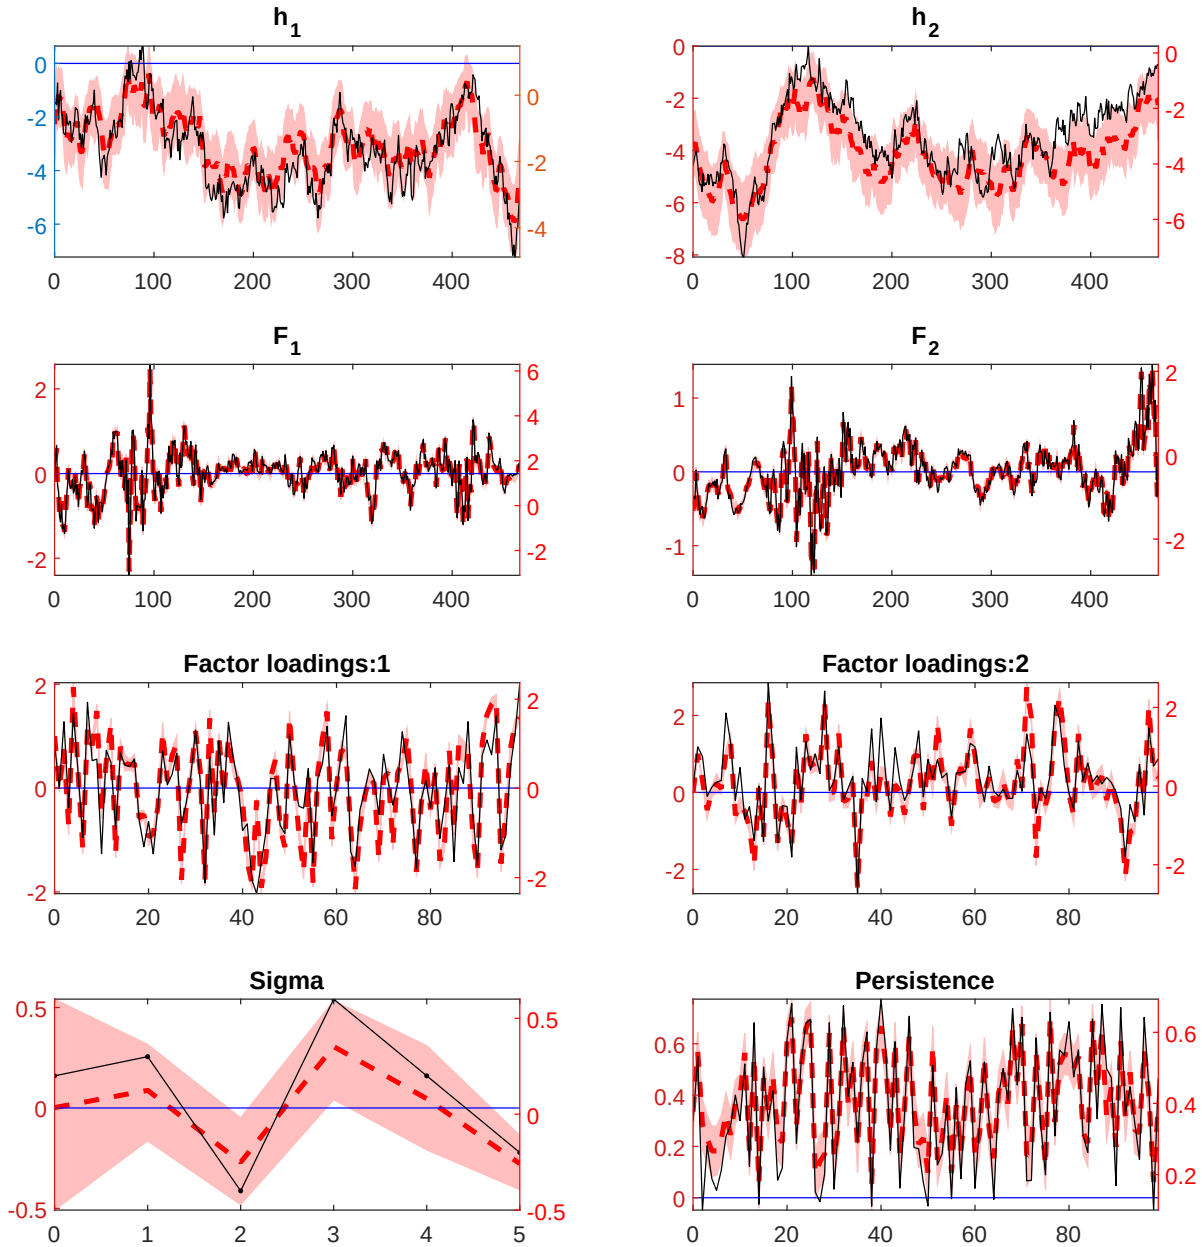
$$\text{with } \text{Corr}(\mathbf{e}_t) = \begin{pmatrix} 1 & 0.2 & 0.3 & -0.4 \\ 0.2 & 1 & 0.6 & 0.2 \\ 0.3 & 0.6 & 1 & -0.2 \\ -0.4 & 0.2 & -0.2 & 1 \end{pmatrix}.$$

We generate $n = 100$ observable series as $X_t = BF_t + v_t$ where B is drawn from $N(0, 1)$, the persistence parameters $\rho_i \sim U(0.1, 0.7)$, and idiosyncratic volatilities evolve as $r_{it} = r_{it-1} + 0.01^{1/2}\epsilon_{it}$.

The model is estimated using 21,000 MCMC draws, discarding the first 1,000 as burn-in and retaining every 10th draw. Figure A.1 compares true parameter values (black) with posterior medians (red dots) and 95% credible intervals (red shaded). The algorithm recovers VAR coefficients, volatility persistence, in-mean effects, and correlation structure accurately. Some scale differences arise because we standardize the data before estimation, but the structural relationships are correctly identified.

This validation confirms that our algorithm can reliably estimate the model's complex interplay of endogenous volatility, mean-volatility correlation, and factor dynamics.

Figure A.1: Estimation on Simulated Data



Note: Estimated Parameters and true values. The two factors are denoted by F_1 and F_2, the log stochastic volatility by h_1 and h_2. The covariance of the error terms of the transition equation for the factors and stochastic volatilities is denoted by *Sigma*, the autoregressive coefficients for the idiosyncratic errors by *persistence* and two columns of the factor loadings matrix by *Factor loadings:1* and *Factor loadings:2*

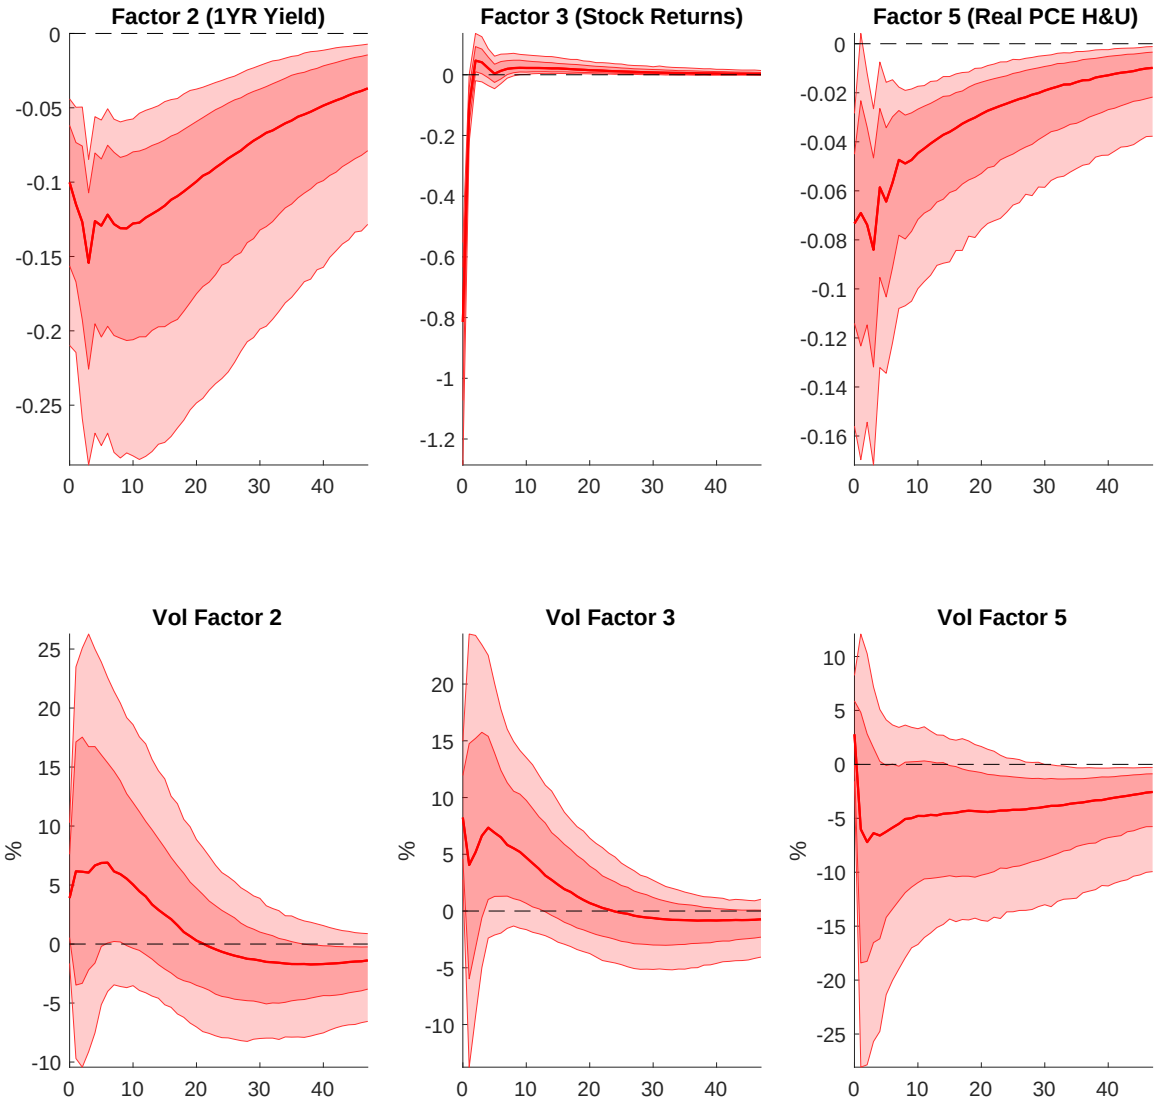
C Effects of Shocks to Financial and Inflation Factors

This section presents the effects of a shock to the first factor (financial) in October 2008 and to the seventh factor (inflation) in December 1980. We assume a Cholesky decomposition with a shock to the financial factor ordered first and shock to the inflation factor ordered second.

Figure A.2 shows the effects of the shock to the financial factor on the second (1 year Treasury yield), third (stock returns), and fifth (housing consumption) factors, which are the responses not shown in the main draft. The shock lowers the Treasury and housing consumption factors persistently, consistent with the effects of a financial shock. The stock return factor sharply declines. The volatilities of the Treasury and stock return factor innovations increase, while the volatility of the housing consumption factor innovations slightly decline.

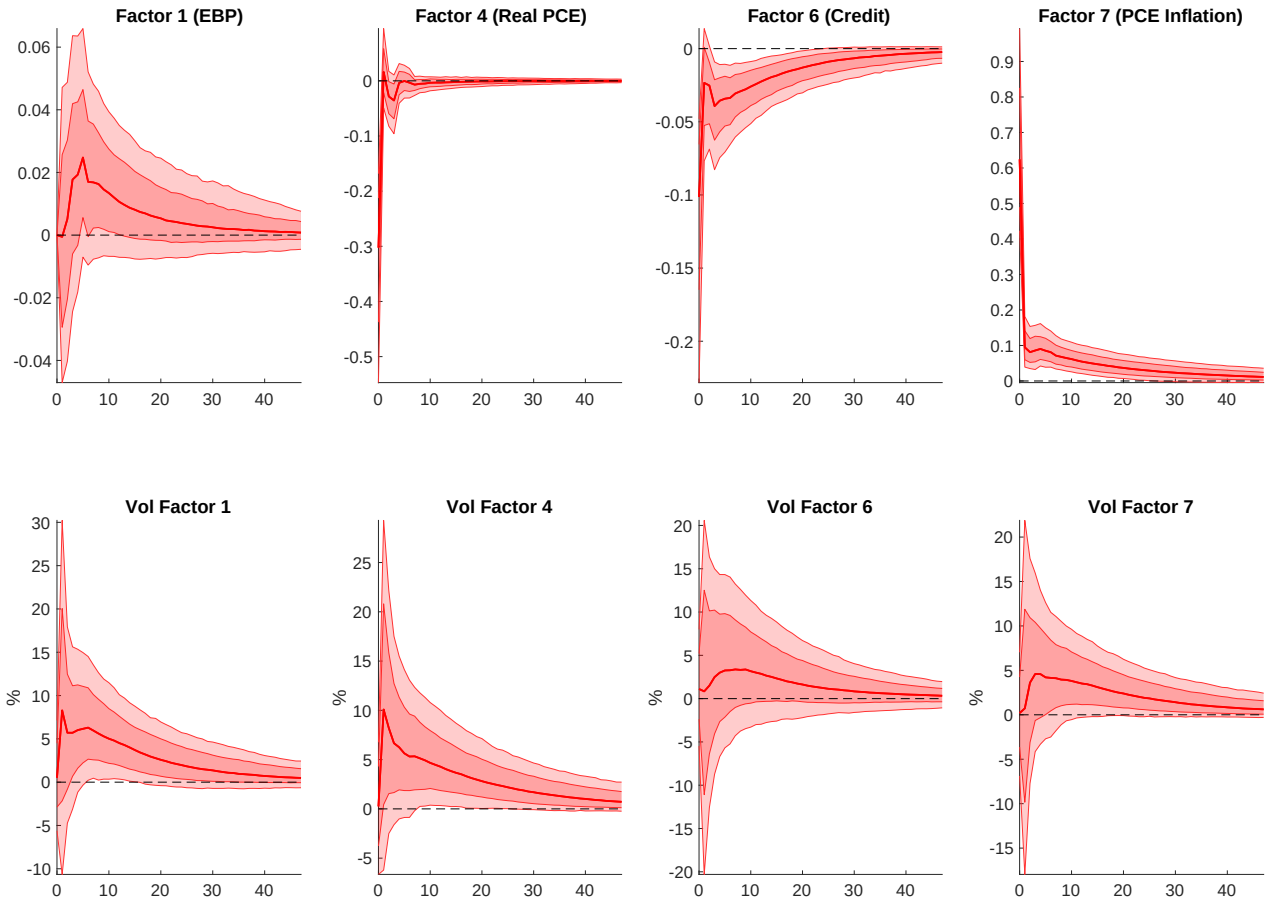
Figures A.3 and A.4 show the responses of the shock to the inflation factor. The inflation factor itself sharply increases before unwinding most of its initial impact and persistently returning to baseline. The consumption and credit factors immediately decline, while the financial factor modestly increases. The Treasury factor increases in response to the increase in inflation, while the stock return and housing consumption factors decline on impact. Volatilities of the innovations across all factors increase following an inflation factor shock, reflecting the destabilizing effects of high inflation.

Figure A.2: Effects of a Shock to the First Factor on Factors and Volatilities in Oct 2008



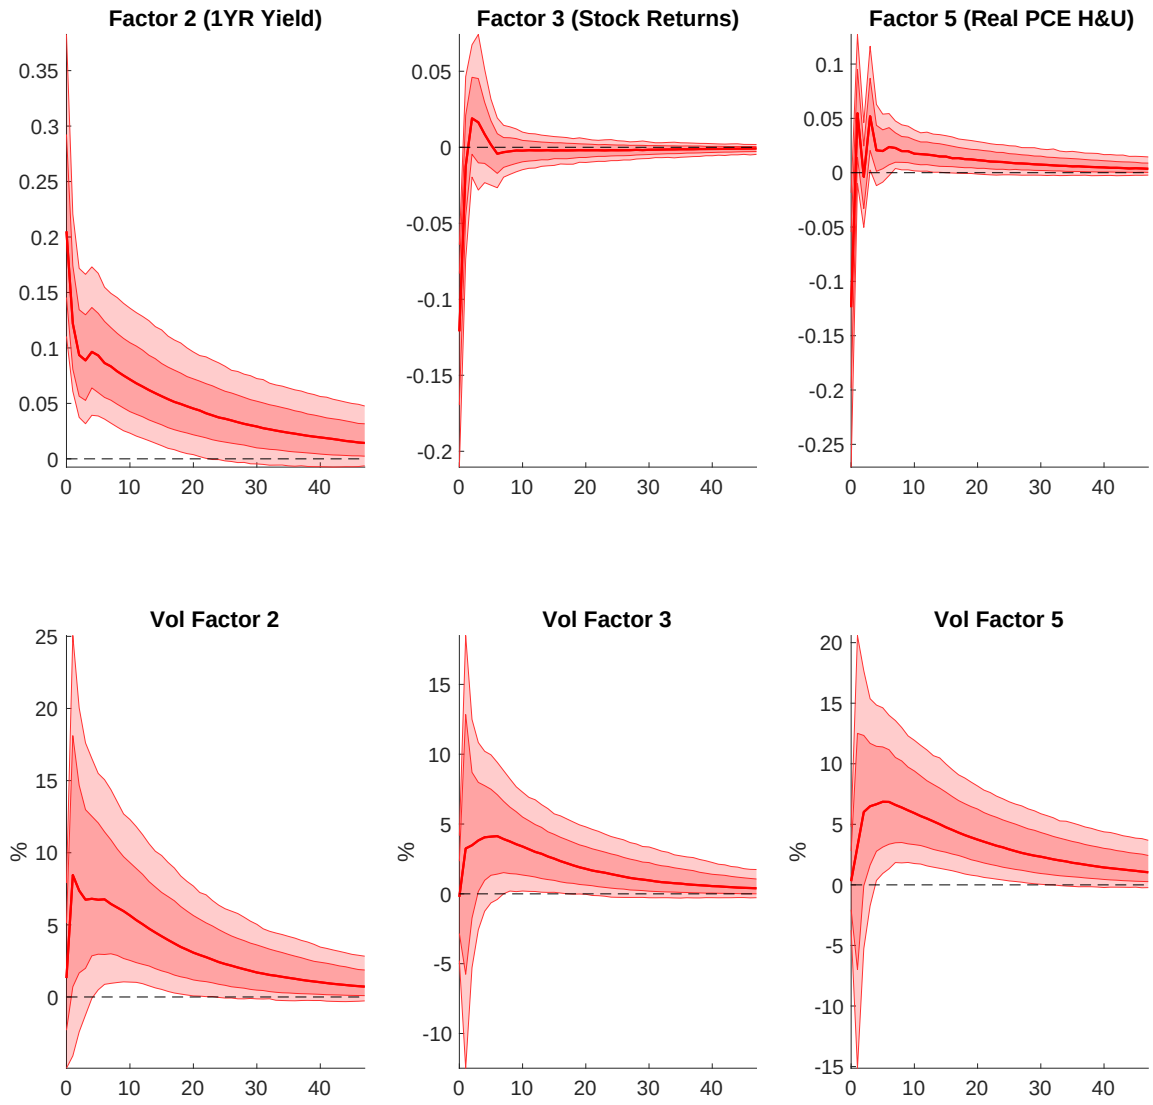
Note: This figure plots impulse response functions of a one standard deviation shock to the first factor in a Cholesky decomposition ordered first given October 2008 conditions. The first column shows the level factor responses and the second column the log volatility responses. The line is the posterior median, the dark shaded area is the 68% credible set and the light shaded area is the 90% credible set.

Figure A.3: Effects of a Shock to the Seventh Factor on Factors and Volatilities in Dec 1980



Note: This figure plots impulse response functions of a one standard deviation shock to the seventh factor in a Cholesky decomposition ordered second given December 1980 conditions. The first column shows the level factor responses and the second column the log volatility responses. The line is the posterior median, the dark shaded area is the 68% credible set and the light shaded area is the 90% credible set.

Figure A.4: Effects of a Shock to the Seventh Factor on Factors and Volatilities in Dec 1980



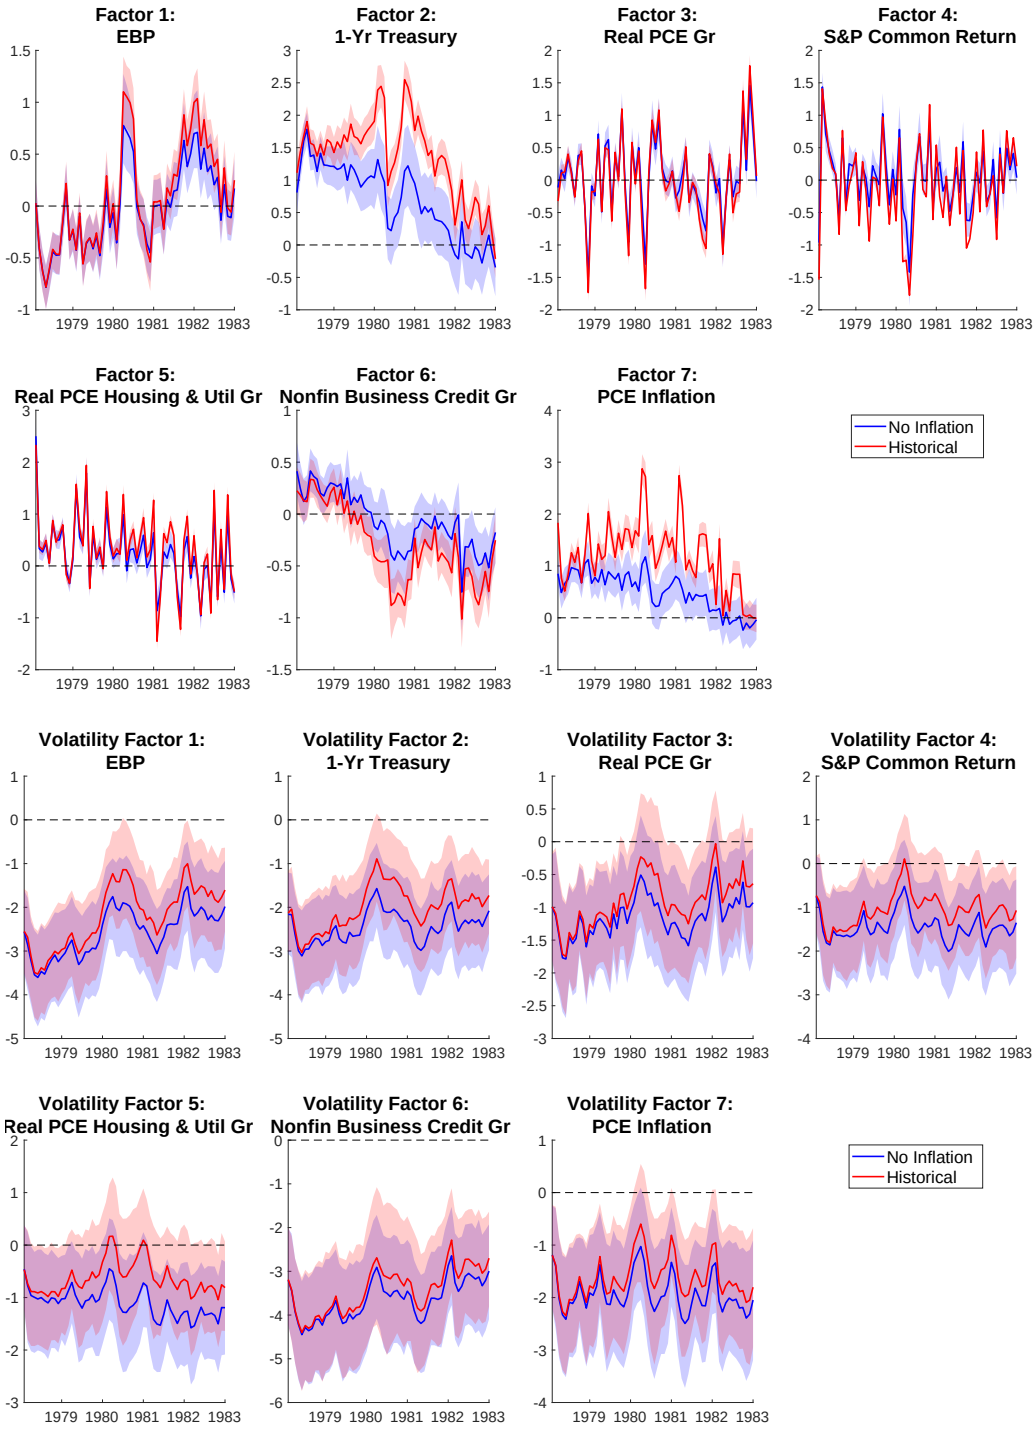
Note: This figure plots impulse response functions of a one standard deviation shock to the seventh factor in a Cholesky decomposition ordered second given December 1980 conditions. The first column shows the level factor responses and the second column the log volatility responses. The line is the posterior median, the dark shaded area is the 68% credible set and the light shaded area is the 90% credible set.

D Details of the Scenario Analysis

In this section, we discuss how we implement the scenario analysis exercise in Section 5.3. The scenario runs from January 1978 - December 1982 for the Great Inflation and from January 2007 - December 2011 for the GFC. We begin with the historical smoothed estimates of the factors and volatilities produced by our estimation, which are shown in the red lines in Figures A.5 and A.6. Then, we form counterfactual paths (blue lines in the figures) by setting either the realized shocks to the financial factor (factor 1) to 0 or shocks to the inflation factor (factor 7) to 0, keeping the other shocks as before. We use a Cholesky ordering with shocks to the financial factor ordered first and shocks to the inflation factor ordered second.

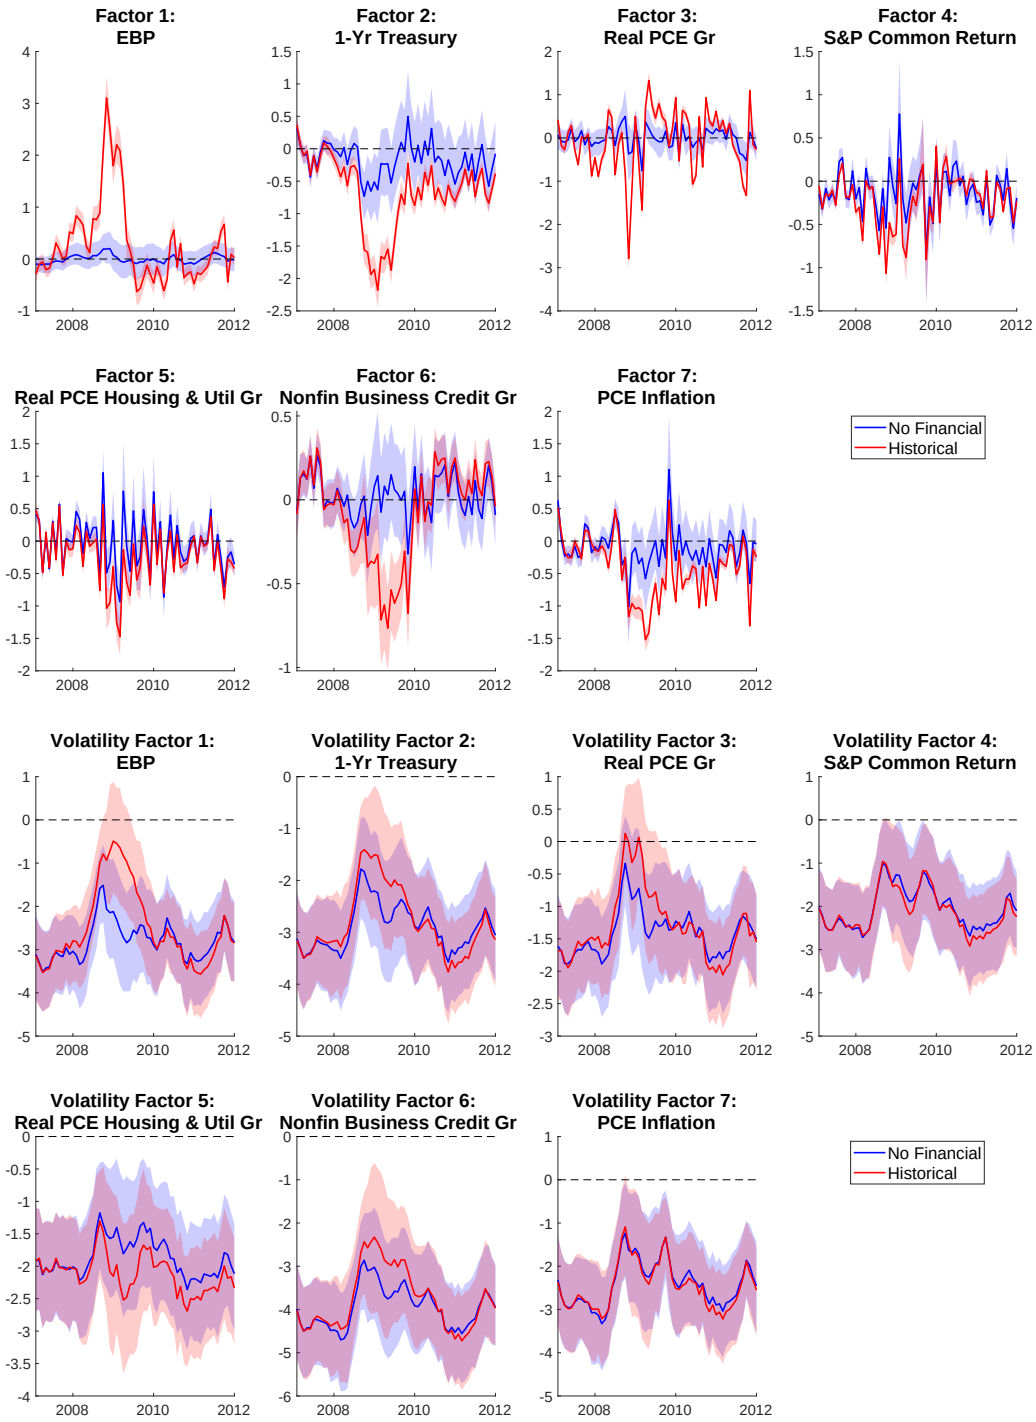
We generate 12−month ahead conditional distributions of the sectoral prices given either the historical or counterfactual states assuming that all shocks are active. We then compute the tail variability statistics shown in Figure 11.

Figure A.5: Great Inflation Scenario: Shutting Off Inflation Shocks



Note: This figure shows the historical evolution of the factors (top) and volatilities (bottom) from January 1978–December 1982. The red lines are the historical estimated evolution of the factors and volatilities along with 90% credible sets. The blue lines are counterfactuals assuming no shocks to the seventh factor ordered second in a Cholesky decomposition occurred over the period.

Figure A.6: GFC Scenario: Shutting Off Financial Shocks



Note: This figure shows the historical evolution of the factors (top) and volatilities (bottom) from January 2007–December 2011. The red lines are the historical estimated evolution of the factors and volatilities along with 90% credible sets. The blue lines are counterfactuals assuming no shocks to the first factor ordered first in a Cholesky decomposition occurred over the period.

E Additional Tables and Figures

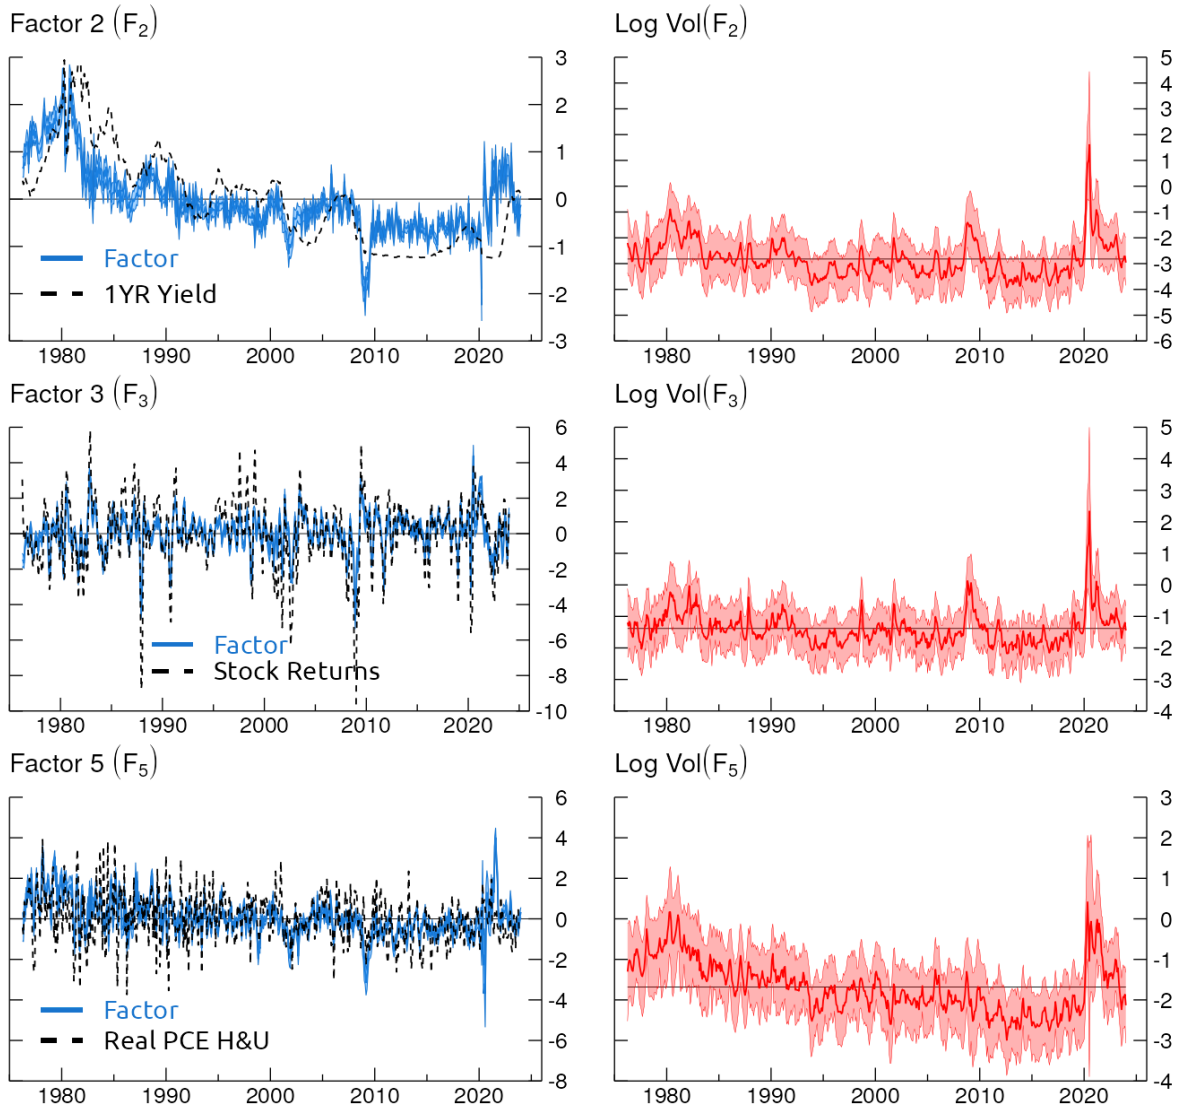
This appendix presents additional tables and figures referenced in the main text: Top 5 Aggregate Indicators by Tail Asymmetry (Table A.3); smoothed estimates and factor volatilities for factors 2, 3, and 5 (Figure A.7); correlation of factors and log volatilities for full sample and subsamples (Figure A.8); and sectoral factor loadings for industrial production (Figure A.9).

Table A.3: Top 5 Aggregate Indicators by Tail Asymmetry

Category	Indicator	Asymmetry ($\text{Log}(\sigma_{P95}/\sigma_{P5})$)
Growth	IP Index	-0.98
	Real Manufacturing and Trade Industries Sales	-0.93
	Initial Claims	-0.90
	New Orders for Durable Goods	-0.78
	All Employees: Total nonfarm	-0.69
Financial	VIX	0.79
	Excess Bond Premium	0.57
	Fed Res 10-year High Quality Corporate Bond Spread	0.31
	1-Year Treasury Rate	0.29
	Baa - 10-Year Treasury	0.26

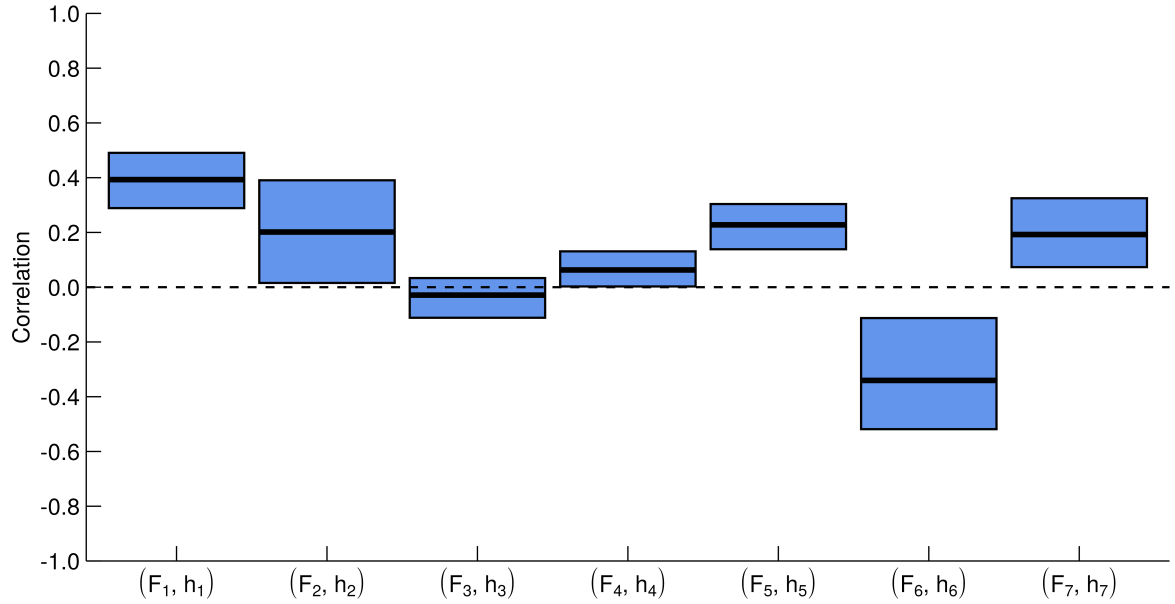
Note: The table shows the top 5 aggregate growth and financial indicators with the most pronounced tail asymmetries, measured as the log ratio of the conditional 95th and 5th percentile standard deviations from July 1976 - June 2019. For growth variables, more negative asymmetry values indicate more downside risk. For financial variables, more positive values indicate more upside (adverse) risk.

Figure A.7: Smoothed Estimates of the Factors and Volatilities

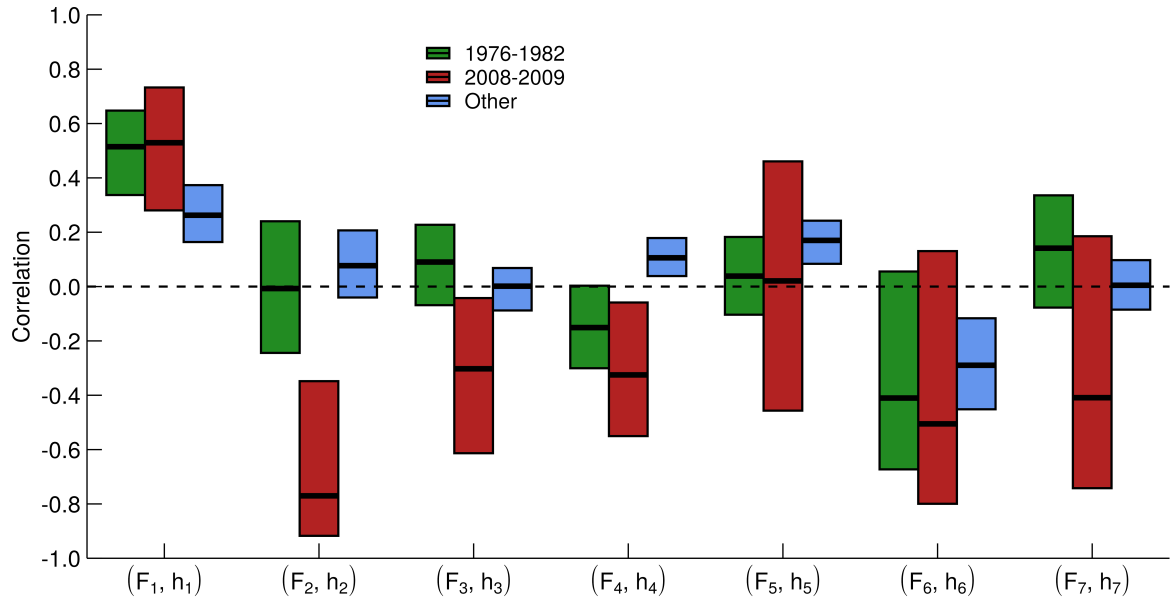


Note: The figure shows posterior median estimates of the level factors (blue lines) and of the log volatilities (red lines). Shaded areas denote the 5th and 95th of the posterior distributions. The dashed lines depict the observable variable associated to the factor for the normalization of matrix B. See Section 2 for details on the normalization.

Figure A.8: Correlation of Factors and Log Volatilities.



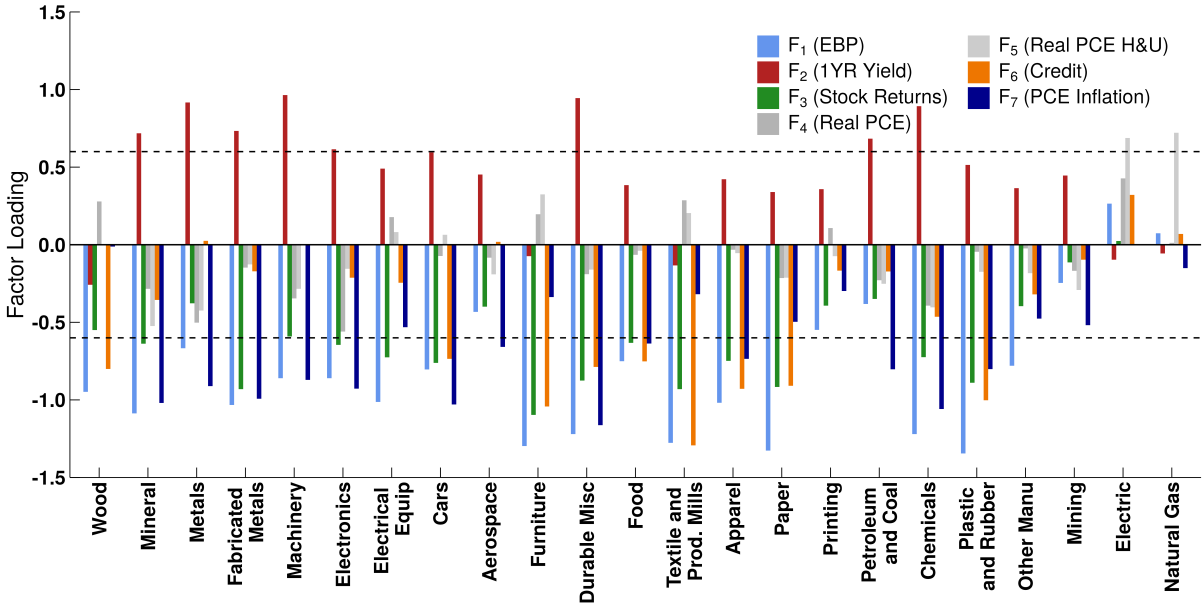
(a) Correlation within Factors: Full Sample.



(b) Correlation within Factors: Subsamples.

Note: Panel (a) shows the pairwise correlation coefficients of the estimated factors F_x and the associated log volatilities h_x. The black lines depict the median correlation, while the boxes depict the min and max elements of the posterior set. Panel (b) shows pairwise correlations across three subsamples: July 1976–Dec 1982, GFC (Dec 2007- June 2009), and remaining months. The boxes depict 70 percent posterior credible sets.

Figure A.9: Sectoral Factor Loadings for Industrial Production



Note: This figure shows the median factor loading estimates for each IP sector. Factors 1, 2, 3, 6, and 7 are colored, representing the five most important factors The two dashed lines denote 0.6 and -0.6.